

Isomorphism Testing for Graphs of Bounded Rank Width

Martin Grohe and Pascal Schweitzer
RWTH Aachen University
{grohe,schweitzer}@informatik.rwth-aachen.de

Abstract

We give an algorithm that, for every fixed k , decides isomorphism of graphs of rank width at most k in polynomial time. As the clique width of a graph is bounded in terms of its rank width, we also obtain a polynomial time isomorphism test for graph classes of bounded clique width.

1 Introduction

Rank width, introduced by Oum and Seymour [26], is a graph invariant that measures how well a graph can be recursively decomposed along “simple separations”. In this sense, it resembles tree width, but it fundamentally differs from tree width in how the “simplicity” of a separation is measured: for rank width, the idea is to take the row rank (over the field \mathbb{F}_2) of the matrix that records the adjacencies between the two parts of a separation, whereas for tree width one simply counts how many vertices the two parts have in common. Rank width is bounded in terms of tree width, but not vice versa. For example, the complete graph K_n has rank width 1 and tree width $n - 1$. This also shows that graphs of bounded rank width are not necessarily sparse (as opposed to graphs of bounded tree width). An interesting aspect of rank width when dealing with problems like graph isomorphism testing (or various problems related to logical definability) that make no real distinction between the edge relation and the “non-edge relation” of a graph is that the rank width of a graph and its complement differ by at most one. Another well-known graph invariant is clique width [6]; it measures how many labels are needed to generate a graph in a certain grammar. Rank width is equivalent to clique width, in the sense that each of the two invariants is bounded in terms of the other [26]. As for bounded tree width, many hard algorithmic problems can be solved in polynomial time (often cubic time) on graph classes of bounded rank width, or equivalently, bounded clique width (e.g. [5, 8, 10, 18]). However, until now it was open whether the isomorphism problem is among them.

We give an algorithm that, for every fixed k , decides isomorphism of graphs of rank width at most k in polynomial time. Many of the best known graph classes where the isomorphism problem is known to be in polynomial time are classes of sparse graphs [16, 9, 25, 24, 28, 2, 14], among them planar graphs, graphs of bounded degree, and graphs of bounded tree width. Less is known for dense graphs; among the known results are polynomial time isomorphism tests for classes with bounded eigenvalue multiplicities [1] and various hereditary graph classes, specifically classes intersection graphs [7, 19], among them interval graphs [23], and classes defined by excluding specific induced subgraphs [3, 20, 32]. Our result substantially extends the realm of hereditary graph classes with a tractable isomorphism problem. While it subsumes several known results [2, 4, 21, 32], for the classes of clique width at most k a polynomial time isomorphism algorithm was only known for the case $k \leq 2$ ([21]).

Technically, we found the isomorphism problem for bounded rank width graphs much harder than anticipated. The overall proof strategy is generic: first compute a canonical decomposition of a graph, or if that is impossible, a canonical family of decompositions with a compact representation, and then use dynamic programming to solve the isomorphism problem. Indeed, this is the strategy taken for bounded tree width graphs in [2, 22]. However, for graphs of bounded

rank width, both steps of this general strategy turned out to be difficult to implement. To compute canonical decompositions, we heavily rely on the general theory of connectivity functions, branch decompositions, and tangles [29, 11], and in particular on computational aspects of the theory recently developed in [15]. Our starting point is an algorithm for canonically decomposing a connectivity function into highly connected regions described by maximal tangles [15]. The technical core of the first part of this paper is a decomposition of these highly connected regions into pieces of bounded width (Lemma 7.1). It has been slightly disturbing to find that even with a canonical decomposition given, the isomorphism problem is still nontrivial and requires a complicated (though elementary) group theoretic machinery. The intuitive reason for this can be explained by a comparison with bounded tree width. In a bounded-width tree decomposition of a graph, we have low order vertex separations of the graph, and after removing the separating vertices (a bounded number) we can deal with the two parts of a separation independently. In a bounded-rank-width decomposition, we have partitions of the graph into two parts such that the adjacency matrix between these parts has low rank. For such a partition, removing a bounded number of vertices shows no effect. Instead, we need to fix a bounded number of rows and columns in the matrix, but even then there is a nontrivial interaction between the two parts, which fortunately we can capture group theoretically.

The paper is organised as follows: after reviewing the necessary background in Section 2, in the short Section 3, we show that all tangles of a connectivity function have “triple covers” of bounded size, providing another technical tool for dealing with tangles (which may be of independent interest). In Section 4, we introduce treelike decompositions of connectivity functions, which may be viewed as compact representations of families of tree decompositions. Sections 5–7 are devoted to a proof of the canonical decomposition theorem (Theorem 7.2). In Section 8, we describe the situation at a single node of our decomposition and its children in matrix form and introduce the notion of partition rank of the matrix to capture the width of the decomposition at this node. Finally, in Sections 9 we develop the group theoretic machinery and give the actual isomorphism algorithm.

Throughout this paper, we often speak of “canonical” constructions. The precise technical meaning depends on the context, but in general a construction (or algorithm) is *canonical* if every isomorphism between its input objects commutes with an isomorphism between the output objects.

2 Connectivity Functions, Tangles, and Branch Decompositions

A *connectivity function* on a finite set A is a symmetric and submodular function $\kappa: 2^A \rightarrow \mathbb{N}$ with $\kappa(\emptyset) = 0$. *Symmetric* means that $\kappa(X) = \kappa(\overline{X})$ for all $X \subseteq A$; here and whenever the ground set A is clear from the context we write \overline{X} to denote $A \setminus X$, the complement of X . *Submodular* means that $\kappa(X) + \kappa(Y) \geq \kappa(X \cap Y) + \kappa(X \cup Y)$ for all $X, Y \subseteq A$. Observe that a symmetric and submodular set function is also *posimodular*, that is, it satisfies $\kappa(X) + \kappa(Y) \geq \kappa(X \setminus Y) + \kappa(Y \setminus X)$ (apply submodularity to X and \overline{Y}).

The only connectivity function that we consider in this paper is the *cut rank* function ρ_G of a graph G . For all subsets $X, Y \subseteq V(G)$, we let $M_{X,Y}$ be the $X \times Y$ -matrix over the 2-element field \mathbb{F}_2 with entries $m_{xy} = 1 \iff xy \in E(G)$. We define $\rho_G: 2^{V(G)} \rightarrow \mathbb{N}$ by letting $\rho_G(X)$ be the row rank of the matrix $M_{X,Y}$ over \mathbb{F}_2 . It is not hard to prove that ρ_G is indeed a connectivity function.

For the rest of this section, let κ be a connectivity function on a finite set A . We often think of a subset $Z \subseteq A$ as a *separation* of A into Z and \overline{Z} and of $\kappa(Z)$ as the *order* of this separation; consequently, we also refer to $\kappa(Z)$ as the *order of Z* . For disjoint sets $X, Y \subseteq A$, an (X, Y) -*separation* is a set $Z \subseteq A$ such that $X \subseteq Z \subseteq \overline{Y}$. Such a separation Z is minimum if its order is minimum, that is, if $\kappa(Z) \leq \kappa(Z')$ for all (X, Y) -separations Z' . It is an easy consequence of the submodularity of κ that there is a unique minimum (X, Y) -separation Z such that $Z \subseteq Z'$ for all other minimum (X, Y) -separations Z' . We call Z the *leftmost minimum*

(X, Y) -separation. There is also a unique *rightmost minimum* (X, Y) -separation, which is easily seen to be the complement of the leftmost minimum (Y, X) -separation.

A κ -tangle of order $k \geq 0$ is a set $\mathcal{T} \subseteq 2^A$ satisfying the following conditions.

- (T.0) $\kappa(X) < k$ for all $X \in \mathcal{T}$,
- (T.1) For all $X \subseteq A$ with $\kappa(X) < k$, either $X \in \mathcal{T}$ or $\overline{X} \in \mathcal{T}$.
- (T.2) $X_1 \cap X_2 \cap X_3 \neq \emptyset$ for all $X_1, X_2, X_3 \in \mathcal{T}$.
- (T.3) \mathcal{T} does not contain any singletons, that is, $\{a\} \notin \mathcal{T}$ for all $a \in A$.

We denote the order of a κ -tangle \mathcal{T} by $\text{ord}(\mathcal{T})$.¹

Let $\mathcal{T}, \mathcal{T}'$ be κ -tangles. If $\mathcal{T}' \subseteq \mathcal{T}$, we say that \mathcal{T} is an *extension* of \mathcal{T}' . The tangles \mathcal{T} and \mathcal{T}' are *incomparable* (we write $\mathcal{T} \perp \mathcal{T}'$) if neither is an extension of the other. The *truncation* of \mathcal{T} to order $k \leq \text{ord}(\mathcal{T})$ is the set $\{X \in \mathcal{T} \mid \kappa(X) < k\}$, which is obviously a tangle of order k . Observe that if \mathcal{T} is an extension of \mathcal{T}' , then $\text{ord}(\mathcal{T}') \leq \text{ord}(\mathcal{T})$, and \mathcal{T}' is the truncation of \mathcal{T} to order $\text{ord}(\mathcal{T}')$.

A κ -tangle \mathcal{T} is *maximal* if there is no κ -tangle $\mathcal{T}' \subseteq \mathcal{T}$ with $\text{ord}(\mathcal{T}') > \text{ord}(\mathcal{T})$. A κ -tangle \mathcal{T} is ℓ -*maximal*, for some $\ell \geq 0$, if either $\text{ord}(\mathcal{T}) = \ell$ or \mathcal{T} is maximal.

A $(\mathcal{T}, \mathcal{T}')$ -separation is a set $Z \subseteq A$ such that $Z \in \mathcal{T}$ and $\overline{Z} \in \mathcal{T}'$. Obviously, if Z is a $(\mathcal{T}, \mathcal{T}')$ -separation then \overline{Z} is a $(\mathcal{T}', \mathcal{T})$ -separation. Observe that there is a $(\mathcal{T}, \mathcal{T}')$ -separation if and only if \mathcal{T} and \mathcal{T}' are incomparable. The *order* of a $(\mathcal{T}, \mathcal{T}')$ -separation Z is $\kappa(Z)$. A $(\mathcal{T}, \mathcal{T}')$ -separation Z is *minimum* if its order is minimum. It can be shown [15] that if $\mathcal{T} \perp \mathcal{T}'$ then there is a unique minimum $(\mathcal{T}, \mathcal{T}')$ -separation Z such that $Z \subseteq Z'$ for all minimum $(\mathcal{T}, \mathcal{T}')$ -separations Z' . We call Z the *leftmost minimum* $(\mathcal{T}, \mathcal{T}')$ -separation. Of course there is also a *rightmost minimum* $(\mathcal{T}, \mathcal{T}')$ -separation, which is the complement of the leftmost minimum $(\mathcal{T}', \mathcal{T})$ -separation.

Now that we have defined (X, Y) -separations for sets X, Y and $(\mathcal{T}, \mathcal{T}')$ -separations for tangles $\mathcal{T}, \mathcal{T}'$, we also need to define combinations of both. For a κ -tangle \mathcal{T} and a set $X \subseteq A$ such that $X \notin \mathcal{T}$, a (\mathcal{T}, X) -separation is a set $Z \in \mathcal{T}$ such that $Z \subseteq \overline{X}$. A (\mathcal{T}, X) -separation is *minimum* if its order is minimum, and again it can be proved that if there is a (\mathcal{T}, X) -separation, then there is a unique *leftmost minimum* (\mathcal{T}, X) -separation and a *rightmost minimum* (\mathcal{T}, X) -separation. Analogously, we define (leftmost, rightmost minimum) (X, \mathcal{T}) -separations.

Lemma 2.1. *Let $\mathcal{T}, \mathcal{T}'$ be κ -tangles and $X, Y \subseteq A$ such that neither $Y \subseteq X$ nor $\overline{Y} \subseteq X$.*

- (1) *If X is a minimum $(\mathcal{T}, \mathcal{T}')$ -separation, then $\kappa(X \cap Y) \leq \kappa(Y)$ or $\kappa(X \cap \overline{Y}) \leq \kappa(Y)$.*
- (2) *If X a rightmost minimum $(\mathcal{T}, \mathcal{T}')$ -separation, then $\kappa(X \cap Y) < \kappa(Y)$ or $\kappa(X \cap \overline{Y}) < \kappa(Y)$.*

Proof. Part (1) is Lemma 4.9 of [15]. We only prove (2). (The proof of (1) is similar.)

Suppose that $\kappa(X \cap Y) \geq \kappa(Y)$ and $\kappa(X \cap \overline{Y}) \geq \kappa(Y)$. By submodularity, $\kappa(X \cup Y) \leq \kappa(X)$ and $\kappa(X \cup \overline{Y}) \leq \kappa(X)$. Then $X \cup Y, X \cup \overline{Y} \in \mathcal{T}$, because both sets are supersets of X . Since $\overline{X} \cap (X \cup Y) \cap (X \cup \overline{Y}) = \emptyset$, either $\overline{X \cup Y} \in \mathcal{T}'$ or $\overline{X \cup \overline{Y}} \in \mathcal{T}'$. If $\overline{X \cup Y} \in \mathcal{T}'$, then $X \cup Y$ is a $(\mathcal{T}, \mathcal{T}')$ -separation whose order is at most the order of X , and as X is a rightmost minimum $(\mathcal{T}, \mathcal{T}')$ -separation, it follows that $X \cup Y \subseteq X$ and thus $Y \subseteq X$. Similarly, if $\overline{X \cup \overline{Y}} \in \mathcal{T}'$, then $\overline{Y} \subseteq X$. \square

The last concept we need to define is that of branch decompositions and branch width of a connectivity function. A *cubic tree* is a tree where every node that is not a leaf has degree 3. An *oriented edge* of a tree T is a pair (s, t) , where $st \in E(T)$. We denote the set of all

¹There is a small technical issue that one needs to be aware of, but that never causes any real problems: if we view tangles as families of sets, then their order is not always well-defined. Indeed, if there is no set X of order $\kappa(X) = k - 1$, then every tangle of order k is equal to its truncation to order $k - 1$. In such a situation, we have to explicitly annotate a tangle with its order, formally viewing a tangle as a pair (\mathcal{T}, k) where $\mathcal{T} \subseteq 2^A$ and $k \geq 0$.

oriented edges of T by $\vec{E}(T)$ and the set of leaves of T by $L(T)$. A *branch decomposition* of κ is a pair (T, ξ) , where T is a cubic tree and $\xi: L(T) \rightarrow A$ is a bijective mapping. For every oriented edge $(s, t) \in \vec{E}(T)$, we let $\tilde{\xi}(s, t) \subseteq A$ be the set of all $\xi(u)$ where u is a leaf in the component of $T - \{st\}$ that contains t (so the oriented edge (s, t) points towards u). Observe that $\tilde{\xi}(s, t) = \tilde{\xi}(t, s)$. We define the *width* of the branch decomposition (T, ξ) to be

$$\text{wd}(T, \xi) := \max\{\kappa(\tilde{\xi}(s, t)) \mid (s, t) \in \vec{E}(s, t)\}.$$

The *branch width* $\text{bw}(\kappa)$ of κ is the minimum of the width of all branch decompositions of κ . The *rank width* of a graph G is defined to be the branch width of the cut rank function ρ_G .

Theorem 2.2 (Duality Theorem [29]). *The branch width of κ is exactly the maximum order of a κ -tangle.*

For disjoint sets $X, Y \subseteq A$ we define $\kappa_{\min}(X, Y) := \min\{\kappa(Z) \mid X \subseteq Z \subseteq \overline{Y}\}$. Note that for all X, Y the two functions $X' \mapsto \kappa_{\min}(X', Y)$ and $Y' \mapsto \kappa_{\min}(X, Y')$ are monotone and submodular.

For sets $Y \subseteq X$, we say that a set Y is *free* in X if $\kappa_{\min}(Y, \overline{X}) = \kappa(X)$ and $|Y| \leq \kappa(X)$. It can be shown that for every $X \subseteq A$ there is a set Y that is free in X [30, 15].

2.1 Computing with Tangles

Algorithms expecting a set function $\kappa: 2^A \rightarrow \mathbb{N}$ as input are given the ground set A as actual input (say, as a list of objects), and they are given an oracle that returns for $X \subseteq A$ the value of $\kappa(X)$. The running time of such algorithms is measured in terms of the size $|A|$ of the ground set. We assume this computation model for all algorithms dealing with abstract connectivity functions κ . Of course, if $\kappa = \rho_G$ is the cut rank function of a graph G , then we assume a standard computation model (without oracles), where the graph G is given as input; we can use G to simulate oracle access to ρ_G .

An important fact underlying most of our algorithms is that, under this model of computation, submodular functions can be efficiently minimised [17, 31].

In [15], we introduced a data structure for representing all tangles of a graph up to a certain order. A *comprehensive tangle data structure* of order k for a connectivity function κ over a set A is a data structure \mathcal{D} with functions $\text{ORDER}_{\mathcal{D}}$, $\text{SIZE}_{\mathcal{D}}$, $\mathcal{T}_{\mathcal{D}}$, $\text{TANGORD}_{\mathcal{D}}$, $\text{TRUNC}_{\mathcal{D}}$, $\text{SEP}_{\mathcal{D}}$, and $\text{FIND}_{\mathcal{D}}$ that provide the following functionalities.

- (1) The function $\text{ORDER}_{\mathcal{D}}()$ returns the fixed integer k .
- (2) For $\ell \in [k]$ the function $\text{SIZE}_{\mathcal{D}}(\ell)$ returns the number of κ -tangles of order at most ℓ . We denote the number of κ -tangles of order at most k by $|\mathcal{D}|$.
- (3) For each $i \in [|\mathcal{D}|]$ the function $\mathcal{T}_{\mathcal{D}}(i, \cdot): 2^A \rightarrow \{0, 1\}$ is a tangle \mathcal{T}_i of order at most k , (i.e., the function call $\mathcal{T}_{\mathcal{D}}(i, X)$ determines whether $X \in \mathcal{T}_i$).

We call i the *index* of the tangle \mathcal{T}_i within the data structure.

- (4) For $i \in [|\mathcal{D}|]$ the call $\text{TANGORD}_{\mathcal{D}}(i)$ returns $\text{ord}(\mathcal{T}_i)$.
- (5) For $i \in [|\mathcal{D}|]$ and $\ell \leq \text{ord}(\mathcal{T}_i)$ the call $\text{TRUNC}_{\mathcal{D}}(i, \ell)$ returns an integer j such that \mathcal{T}_j is the truncation of \mathcal{T}_i to order ℓ . If $\ell > \text{ord}(\mathcal{T}_i)$ the function returns i .
- (6) For distinct $i, j \in [|\mathcal{D}|]$ the call $\text{SEP}_{\mathcal{D}}(i, j)$ outputs a set $X \subseteq A$ such that X is the leftmost minimum $(\mathcal{T}_i, \mathcal{T}_j)$ -separation or states that no such set exists (in which case one of the tangles is a truncation of the other).

- (7) Given $\ell \in \{0, \dots, k\}$ and a tangle \mathcal{T}' of order ℓ (via a membership oracle) the function $\text{FIND}_{\mathcal{D}}(\ell, \mathcal{T}')$, returns the index of \mathcal{T}' , that is, the unique integer $i \in [|\mathcal{D}|]$ such that $\text{ord}(\mathcal{T}_i) = \ell$ and $\mathcal{T}' = \mathcal{T}_i$.

Theorem 2.3 ([15]). *For every constant k there is a polynomial time algorithm that, given oracle access to a connectivity function κ , computes an efficient comprehensive tangle data structure of order k .*

Using a comprehensive tangle data structure, we can design polynomial time algorithms for other computational problems related to tangles.

Lemma 2.4. *Let $k \geq 0$.*

- (1) *There is a polynomial time algorithm that, given a set $X \subseteq A$ and a tangle \mathcal{T} of order k (via its index in a comprehensive tangle data structure), computes the leftmost minimum (\mathcal{T}, X) -separation if it exists or reports that there is no (\mathcal{T}, X) -separation.*
- (2) *There is a polynomial time algorithm that, given a tangle \mathcal{T} of order k (via its index in a comprehensive tangle data structure), computes a list of all inclusionwise minimal elements of \mathcal{T} .*

Proof. Assertion (1) follows from Lemma 2.20 of [15].

To prove (2), we claim that a set $X \subseteq A$ is an inclusionwise minimal element of a tangle \mathcal{T} of order k if and only if the following two conditions are satisfied.

- (i) There is a set $Y \subseteq \overline{X}$ of size $|Y| \leq k$ such that X is the leftmost minimum (\mathcal{T}, Y) -separation.
- (ii) There is no set $Z \subseteq A$ of size $|Z| \leq k$ such that the leftmost minimum (\mathcal{T}, Z) -separation is a proper subset of X .

To see this, we simply observe that if X is an inclusionwise minimal element of \mathcal{T} , then it trivially satisfies (ii), and it satisfies (i), because we can let Y be a set that is free in \overline{X} . Conversely, if X satisfies (i) then it is an element of \mathcal{T} , and (ii) makes sure that it is inclusionwise minimal.

There are at most $\binom{|A|}{k}$ sets X satisfying (i), and using (1) we can list these in polynomial time. Then, using (1) again, for each of these sets we can check whether they satisfy (ii). \square

2.2 Contractions

Contractions give a way to construct new connectivity functions from given ones. To define a *contraction*, we take one or several disjoint subsets of the ground set and “contract” these sets to single points. In the new decomposition, these new points represent the sets of the original decomposition

For the formal treatment, let κ be a connectivity function on a set A .

Let $C_1, \dots, C_m \subseteq A$ be mutually disjoint subsets of A . Let $B := A \setminus (C_1 \cup \dots \cup C_m)$, and let c_1, \dots, c_m be fresh elements (mutually distinct, and distinct from all elements of B). We define

$$A \downarrow_{C_1, \dots, C_m} := B \cup \{c_1, \dots, c_m\}.$$

To simplify the notation, here and in the following we omit the index C_1, \dots, C_m if the sets C_i are clear from the context. For every subset $X \subseteq A \downarrow$, we define its *expansion* to be the set

$$X \uparrow := X \uparrow_{C_1, \dots, C_m} := (X \cap B) \cup \bigcup_{\substack{i \in [m] \\ c_i \in X}} C_i.$$

The C_1, \dots, C_m -contraction of κ is the function $\kappa \downarrow$, or $\kappa \downarrow_{C_1, \dots, C_m}$, on $2^{A \downarrow}$ defined by

$$\kappa \downarrow(X) := \kappa(X \uparrow).$$

It is easy to verify that $\kappa \downarrow$ is indeed a connectivity function.

Remark 2.5. A different view on contractions is to maintain the ground set, but define the connectivity function on a sublattice of the power set lattice. That is, not all separations of the ground set get an order, but only some of them.

Formally, we let $\mathcal{L} := \mathcal{L}(A \downarrow C_1, \dots, C_m)$ be the sublattice of $\mathcal{P}(A) := (2^A, \cap, \cup)$ consisting of all sets $X \subseteq A$ such that for all $i \in [m]$ either $C_i \subseteq X$ or $C_i \subseteq \overline{X}$. Obviously, \mathcal{L} is closed under intersection and union and thus indeed a sublattice. Observe that every $X \in \mathcal{L}$ has a natural contraction

$$X \downarrow := (X \cap B) \cup \{c_i \mid i \in [m] \text{ with } C_i \subseteq X\},$$

and we have $X \downarrow \uparrow = X$. As we also have $X' \uparrow \in \mathcal{L}$ for all $X' \subseteq A \downarrow$, the contraction mapping is a bijection between $\mathcal{L}(A \downarrow C_1, \dots, C_m)$ and $A \downarrow$. It follows immediately from the definition of $\kappa \downarrow$ that for all $X \in \mathcal{L}(A \downarrow C_1, \dots, C_m)$ we have

$$\kappa(X) = \kappa \downarrow (X \downarrow).$$

Thus the contraction mapping is an isomorphism from the *connectivity system* $(\mathcal{L}, \kappa|_{\mathcal{L}})$, where $\kappa|_{\mathcal{L}}$ denotes the restriction of κ to \mathcal{L} , and the connectivity system $(2^{A \downarrow}, \kappa \downarrow)$.

The view of a contraction of κ as a restriction to a sublattice will be useful when dealing with contractions of the cut-rank function of a graph in Section 6.2. \lrcorner

Let \mathcal{T} be a κ -tangle of order k . We define

$$\mathcal{T} \downarrow := \mathcal{T} \downarrow_{C_1, \dots, C_m} := \{X \subseteq A \downarrow \mid X \uparrow \in \mathcal{T}\}.$$

Note that $\mathcal{T} \downarrow$ is not necessarily a $\kappa \downarrow$ -tangle: if $C_i \in \mathcal{T}$ for some $i \in [m]$, then $\{c_i\} \in \mathcal{T} \downarrow$, and thus $\mathcal{T} \downarrow$ violates (T.3). However, it is straightforward to verify that $\mathcal{T} \downarrow$ is a $\kappa \downarrow$ -tangle (of the same order k) if and only if $C_1, \dots, C_m \notin \mathcal{T}$.

3 Triple Covers

A *cover* of a κ -tangle \mathcal{T} is a set $C \subseteq A$ such that $C \cap Y \neq \emptyset$ for all $Y \in \mathcal{T}$. It is not hard to prove that every κ -tangle of order k has a cover of size at most k .

A *triple cover* of a \mathcal{T} is a set $Q \subseteq A$ such that $Q \cap Y_1 \cap Y_2 \cap Y_3 \neq \emptyset$ for all $Y_1, Y_2, Y_3 \in \mathcal{T}$. We shall prove that every tangle of order k has a triple cover of size bounded in terms of k .

Observe that we can test in polynomial time whether a given set Q is a triple cover for a κ -tangle \mathcal{T} , given by its index in a comprehensive tangle data structure: using the data structure, we produce a list of all inclusionwise minimal elements of \mathcal{T} , and then we check if any three of them have a nonempty intersection with Q .

Let $\theta : \mathbb{N} \rightarrow \mathbb{N}$ be defined by $\theta(0) := 0$ and

$$\theta(i+1) := \theta(i) + 3^{\theta(i)}.$$

Lemma 3.1. *Let \mathcal{T} be a κ -tangle of order k . Then \mathcal{T} has a triple cover of size at most $\theta(3k-2)$.*

Proof. By induction on $i \geq 0$ we construct sets Q_i such that for all $Y_1, Y_2, Y_3 \in \mathcal{T}$, if $Q_i \cap Y_1 \cap Y_2 \cap Y_3 = \emptyset$ then $\kappa(Y_1) + \kappa(Y_2) + \kappa(Y_3) \geq i$. Then $Q := Q_{3k-2}$ is a triple cover of \mathcal{T} .

We let $Q_0 := \emptyset$.

For the inductive step $i \rightarrow i+1$, suppose that Q_i is defined. Let \mathcal{X}_i be the set of all partitions (X_1, X_2, X_3) of Q_i into three possibly empty sets. Then $|\mathcal{X}_i| \leq 3^{\theta(i)}$.

For $X = (X_1, X_2, X_3) \in \mathcal{X}_i$ we shall define an element $y_X \in A$ such that for all $Y_1, Y_2, Y_3 \in \mathcal{T}$, if $Q_i \cap Y_j \subseteq X_j$ for $j = 1, 2, 3$ then either $y_X \in Y_1 \cap Y_2 \cap Y_3$ or $\kappa(Y_1) + \kappa(Y_2) + \kappa(Y_3) \geq i+1$. Then we let

$$Q_{i+1} := Q_i \cup \{y_X \mid X \in \mathcal{X}_i\}.$$

Clearly, $|Q_{i+1}| \leq |Q_i| + |\mathcal{X}_i| \leq \theta(i+1)$, and if $Y_1, Y_2, Y_3 \in \mathcal{T}$ with $Q_{i+1} \cap Y_1 \cap Y_2 \cap Y_3 = \emptyset$, then $\kappa(Y_1) + \kappa(Y_2) + \kappa(Y_3) \geq i+1$, because otherwise

$$y_{((Q_i \cap Y_1), (Q_i \cap Y_2), (Q_i \cap Y_3))} \in Y_1 \cap Y_2 \cap Y_3.$$

Observe that for every $X \subseteq Q_i$, either there is no $Y \in \mathcal{T}$ such that $Q_i \cap Y \subseteq X$ or there is a (unique) $Y \in \mathcal{T}$ such that

- (i) $Q_i \cap Y \subseteq X$;
- (ii) subject to (i), $\kappa(Y)$ is minimum;
- (iii) $Y \subseteq Y'$ for all $Y' \in \mathcal{T}$ satisfying (i) and (ii).

This can be proved by a standard submodularity argument.

Now let $X = (X_1, X_2, X_3) \in \mathcal{X}_i$. If for some $j \in \{1, 2, 3\}$, there is no $Y_j \in \mathcal{T}$ such that $Q_i \cap Y_j \subseteq X_j$, then there is nothing to do, and we can choose y_X arbitrarily. Otherwise, for $j = 1, 2, 3$ we let Y_j be the unique set in \mathcal{T} satisfying (i)–(iii) with respect to X_j . Then $Q_i \cap Y_1 \cap Y_2 \cap Y_3 \subseteq X_1 \cap X_2 \cap X_3 = \emptyset$ and thus, by the induction hypothesis, $\kappa(Y_1) + \kappa(Y_2) + \kappa(Y_3) \geq i$.

Let $y = y_X \in Y_1 \cap Y_2 \cap Y_3$; such a y exists by (T.2). Let $Z_1, Z_2, Z_3 \in \mathcal{T}$ such that $Q_i \cap Z_j \subseteq X_j$ and $y \notin Z_1 \cap Z_2 \cap Z_3$; if no such Z_j exist there is nothing to prove. We claim that

$$\kappa(Z_1) + \kappa(Z_2) + \kappa(Z_3) > \kappa(Y_1) + \kappa(Y_2) + \kappa(Y_3) \geq i. \quad (3.A)$$

We observe first that for $j = 1, 2, 3$ we have $\kappa(Z_j) \geq \kappa(Y_j)$ by (ii), and if $\kappa(Z_j) = \kappa(Y_j)$ then $Y_j \subseteq Z_j$ by (iii).

Without loss of generality we may assume that $y \notin Z_1$. Then $Y_1 \not\subseteq Z_1$ and thus $\kappa(Z_1) > \kappa(Y_1)$. This proves (3.A). \square

4 Treelike Decompositions

In a directed graph D , by $N_+^D(t)$ or just $N_+(t)$ if D is clear from the context, we denote the set out-neighbours of a node t . By \leq^D or just \leq we denote the reflexive transitive closure of $E(D)$, which is a partial order if D is acyclic. A *directed tree* is a directed graph T where for all nodes t the set $\{s \mid s \leq t\}$ is linearly ordered by \leq .

Let A be a set. A *directed decomposition* of A is a pair (D, γ) , where D is a directed graph and $\gamma : V(D) \rightarrow 2^A$. If κ is a connectivity function on A , we also say that (D, γ) is a directed decomposition of κ . For every node $t \in V(D)$, we let

$$\beta(t) := \gamma(t) \setminus \bigcup_{u \in N_+^D(t)} \gamma(u). \quad (4.A)$$

We call $\beta(t)$ the *bag* and $\gamma(t)$ the *cone* at t . We always denote the bag function of a directed decomposition (D, γ) by β , and we use implicit naming conventions by which, for example, we denote the bag function of (D', γ') by β' .

A directed decomposition (D, γ) of A is *treelike*, or a *treelike decomposition*, if it satisfies the following axioms.

(TL.1) D is a acyclic.

(TL.2) For all $(t, u) \in E(D)$,

$$\gamma(t) \supseteq \gamma(u).$$

(TL.3) For all $t \in V(T)$ and $u_1, u_2 \in N_+^D(t)$,

$$\gamma(u_1) = \gamma(u_2) \quad \text{or} \quad \gamma(u_1) \cap \gamma(u_2) = \emptyset.$$

(TL.4) There is a $t \in V(D)$ such that $\gamma(t) = A$.

If (D, γ) only satisfies (TL.1)–(TL.3), we call it a *partial treelike decomposition*. The treelike decompositions of connectivity functions introduced here are adaptations of treelike decompositions of graphs introduced in [12, 13].

In the following, let (D, γ) be a partial treelike decomposition of A . Observe that for all $t \in V(D)$,

$$\gamma(t) = \bigcup_{u \succeq t} \beta(u). \quad (4.B)$$

(D, γ) is a (partial) *directed tree decomposition*² if D is a directed tree and for all $t \in V(D)$ and all distinct $u_1, u_2 \in N_+^D(t)$,

$$\gamma(u_1) \cap \gamma(u_2) = \emptyset.$$

Observe that (D, γ) is a directed tree decomposition if and only if D is a directed tree and the bags $\beta(t)$ for $t \in V(D)$ are mutually disjoint and have union A (that is, they form a partition of A with possibly empty parts).

Now assume that κ is a connectivity function on A and (D, γ) a (partial) treelike decomposition of κ . The *width* of a node $t \in V(D)$ in (D, γ) is

$$\text{wd}(D, \gamma, t) := \max_{\substack{X \subseteq \beta(t) \\ U \subseteq N_+^D(t)}} \kappa \left(X \cup \bigcup_{u \in U} \gamma(u) \right).$$

The *width* $\text{wd}(D, \gamma)$ of the decomposition is the maximum of the widths of its nodes.

It is sometimes convenient to normalise treelike decompositions. The *roots* of a directed acyclic graph D are the nodes r of in-degree 0. The *leaves* of a directed acyclic graph D are the nodes t of out-degree 0; we denote the set of all leaves of D by $L(D)$. All non-leaf nodes are called *inner nodes*. We say that a treelike decomposition (D, γ) of κ is *normal* if (in addition to the axioms (TL.1)–(TL.4)) it satisfies the following conditions.

(NTL.1) For all inner nodes $t \in V(D) \setminus L(D)$ it holds that $\beta(t) = \emptyset$.

(NTL.2) For all leaves $t \in L(D)$ it holds that $|\beta(t)| = 1$.

(NTL.3) For all nodes $t \in V(D)$, either $\gamma(u_1) = \gamma(u_2)$ for all $u_1, u_2 \in N_+(t)$ or $\gamma(u_1) \cap \gamma(u_2) = \emptyset$ for all distinct $u_1, u_2 \in N_+(t)$.

(NTL.4) D has a unique root.

Lemma 4.1. *Let (D, γ) be a treelike decomposition of κ . Then there is a normal treelike decomposition (D', γ') of κ such that $\text{wd}(D, \gamma) = \text{wd}(D', \gamma')$. If D is a tree, then D' is a tree as well, and if (D, γ) is a tree decomposition then (D', γ') is a tree decomposition as well.*

Furthermore, the construction of (D', γ') from (D, γ) is canonical and can be carried out in polynomial time.

Proof. To satisfy (NTL.1) and (NTL.2), we extend our decomposition as follows. Let t be an inner node with $\beta(t) \neq \emptyset$ or a leaf $t \in L(D)$ with $|\beta(t)| > 1$. For each $x \in \beta(t)$, we add a fresh node t_x and an edge from t to t_x . We leave $\gamma(v)$ unchanged for all old nodes v and set $\gamma(t_x) := \{x\}$. We obtain a new treelike decomposition that satisfies (NTL.1) and (NTL.2). This new decomposition has the same width as the old one.

To satisfy (NTL.3), we modify our decomposition as follows. For every node $t \in V(D)$ that has distinct children $u_1, u_2, u_3 \in N_+(t)$ such that $\gamma(u_1) = \gamma(u_2) \neq \gamma(u_3)$, we partition $N_+(t)$ into sets U_1, \dots, U_m such that all nodes in U_i have the same cone and the nodes in distinct U_i, U_j have disjoint cones. We delete all edges from t to its children. Then we add m fresh

²Deviating from previous work [11, 15], we view the trees in tree decompositions as being directed.

nodes t_1, \dots, t_m and edges from t to t_i and from t_i to all nodes in U_i . We leave $\gamma(v)$ unchanged for all old nodes v and set $\gamma(t_i) := \gamma(u_i)$ for some (and hence all) $u_i \in U_i$. We obtain a new treelike decomposition that satisfies (NTL.3). This new decomposition has the same width as the old one, and it also leaves (NTL.1) and (NTL.2) intact.

Suppose now that (D, γ) satisfies (NTL.1)–(NTL.3); it remains to satisfy (NTL.4). We repeatedly remove roots t with $\gamma(t) \neq A$ until we are left with a graph D' where all roots we have $\gamma(t) = A$. Then we add a new node r and edges from r to all roots of D' ; the resulting graph D'' has r as its only root. We define $\gamma'' : V(D'') \rightarrow 2^A$ by $\gamma''(r) := A$ and $\gamma''(t) := \gamma(t)$ for all $t \in V(D') \subseteq V(D)$. It is easy to see that (D'', γ'') satisfies (TL.1)–(TL.4) and thus is a treelike decomposition. It satisfies (NTL.4) by construction. We have $\beta''(r) = \emptyset$ and $\beta''(t) = \beta(t)$ for all $t \in V(D')$. Thus the decomposition (D'', γ'') satisfies (NTL.1) and (NTL.2). It also satisfies (NTL.3), because all children of the root r have the same cone A and all nodes $t \in V(D')$ have the same children as in D . Finally, (D'', γ'') has the same width as (D, γ) , because the width at the root r is 0 and the width at all other nodes t is the same as in (D, γ) . \square

The following lemma shows that our definition of width is fairly robust.

Lemma 4.2. *For every $k \in \mathbb{N}$, the following statements are equivalent.*

- (i) $\text{bw}(\kappa) \leq k$.
- (ii) κ has a directed tree decomposition of width at most k .
- (iii) κ has a treelike decomposition of width at most k .

Proof. To prove (i) \implies (ii), let (T, ξ) be a branch decomposition of κ . If $E(T) = \emptyset$, then $|A| = 1$, and the claim is trivial. So we assume $E(T) \neq \emptyset$. Let $e_0 = s_0 t_0 \in E(T)$. We define a directed tree T' by subdividing the edge e_0 , making the newly inserted node, say, r , the root of T' and directing all edges away from r . We define $\gamma : V(T') \rightarrow 2^A$ by:

- $\gamma(r) := A$;
- $\gamma(s_0) := \tilde{\xi}(t_0, s_0)$ and $\gamma(t_0) := \tilde{\xi}(s_0, t_0)$;
- $\gamma(t) := \tilde{\xi}(s, t)$ for every node $t \in V(T') \setminus \{r, s_0, t_0\}$ with parent s .

Observe every internal node $t \in V(T') \setminus L(T')$ has precisely two children, because the tree T is cubic.

Clearly, (T', γ) is a directed tree decomposition of κ .

Claim 1. $\text{wd}(T', \gamma) \leq \text{wd}(T, \xi)$.

Proof. Let $t \in V(T')$; we shall prove that $\text{wd}(T', \gamma, t) \leq \text{wd}(T, \xi)$.

Suppose first that $t \in L(T)$. Then $\gamma(t) = \beta(t) = \{\xi(t)\}$, and thus

$$\text{wd}(T, \gamma, t) = \max\{\kappa(\emptyset), \kappa(\{\xi(t)\})\} = \kappa(\{\xi(t)\}) \leq \text{wd}(T, \xi).$$

Suppose next that t has parent $s \neq r$ and children u_1, u_2 . Then $\overline{\gamma(t)} = \tilde{\xi}(t, s)$, $\gamma(u_1) = \tilde{\xi}(t, u_1)$, and $\gamma(u_2) = \tilde{\xi}(t, u_2)$. As the union of these three sets is A , we have $\gamma(t) = \gamma(u_1) \cup \gamma(u_2)$ and thus $\beta(t) = \emptyset$. Thus

$$\begin{aligned} \text{wd}(T', \gamma, t) &= \max\{\kappa(\emptyset), \kappa(\gamma(u_1)), \kappa(\gamma(u_2)), \kappa(\gamma(t))\} \\ &= \max\{0, \kappa(\tilde{\xi}(t, u_1)), \kappa(\tilde{\xi}(t, u_2)), \kappa(\tilde{\xi}(s, t))\} \leq \text{wd}(T, \xi). \end{aligned}$$

If $t = s_0$ or t_0 , we can argue completely analogously, using t_0 or s_0 , respectively, in place of the parent s .

Finally, suppose that $t = r$ is the root. Then $N_+(t) = \{s_0, t_0\}$, and thus

$$\beta(t) = \gamma(t) \setminus (\gamma(s_0) \cup \gamma(t_0)) = A \setminus (\tilde{\xi}(t_0, s_0) \cup \tilde{\xi}(s_0, t_0)) = \emptyset,$$

because $\xi(s_0, t_0) = \overline{\xi(t_0, s_0)}$. It follows that

$$\text{wd}(T, \gamma, t) = \max\{\kappa(\emptyset), \kappa(\tilde{\xi}(t_0, s_0)), \kappa(\tilde{\xi}(s_0, t_0)), \kappa(A)\} \leq \text{wd}(T, \xi). \quad \lrcorner$$

To prove (ii) \implies (i), let (T, γ) be a directed tree decomposition of κ . Without loss of generality we may assume that it is normal. If T is a binary tree (where each internal node has exactly two children), we can turn the decomposition into a branch decomposition simply by inverting the construction in the proof of (i) \implies (ii).

It remains to turn the tree T into a binary tree. Suppose that $t \in V(T)$ is a node with at least three children. Let U_1, U_2 be a partition of $U := N_+(t)$. We modify the tree by deleting all edges from t to its children, inserting two new children u_1, u_2 for t , and making all nodes in U_1 children of u_1 and all nodes in U_2 children of u_2 . Let T' be the resulting tree. We define $\gamma' : V(T') \rightarrow 2^A$ by letting $\gamma'(u_i) := \bigcup_{u \in U_i} \gamma(u)$ for $i = 1, 2$ and $\gamma'(x) := \gamma(x)$ for all $x \in V(T') \setminus \{u_1, u_2\} = V(T)$. Observe that $\text{wd}(T', \gamma', u_i) \leq \text{wd}(T, \gamma, t)$ for $i = 1, 2$ and

$$\text{wd}(T', \gamma', t) = \max\{\kappa(\emptyset), \kappa(\gamma(u_1)), \kappa(\gamma(u_2)), \kappa(\gamma(t))\} \leq \text{wd}(T, \gamma, t)$$

and $\text{wd}(T', \gamma', x) \leq \text{wd}(T, \gamma, x)$ for all $x \in V(T') \setminus \{u_1, u_2, t\}$. Thus $\text{wd}(T', \gamma') \leq \text{wd}(T, \gamma)$. We repeat this construction until the tree is binary.

The implication (ii) \implies (iii) is trivial, because every directed tree decomposition is a treelike decomposition. Thus it remains to prove (iii) \implies (ii). Let (D, γ) be a treelike decomposition of κ . By repeatedly duplicating subtrees, starting from the leafs, we can turn D into a forest F (which may be exponentially larger than D). We can define a function $\gamma' : V(F) \rightarrow 2^A$ accordingly and obtain a treelike decomposition (F, γ') of the same width that is based on the forest F instead of the directed acyclic graph D . We pick a root node r of some tree T of F with $\gamma'(r) = A$. Then we prune the tree as follows: whenever we have a node t with a family $U \subseteq N_+(T)$ of children that all have the same $\gamma(u)$, we delete all but one of these children and the whole subtrees attached to them. The result is a tree decomposition. \square

Remark 4.3. Note that the constructions in the proof of Lemma 4.2, turning treelike decomposition into a tree decomposition and a tree decomposition into a branch decomposition, are not canonical. In fact, it is not difficult to see that there are no width-preserving canonical constructions for these tasks; this is why we introduced treelike decompositions and tree decompositions in the first place.

Now let \mathfrak{T} be a family of mutually incomparable κ -tangles. A *directed tree decomposition* for \mathfrak{T} is a triple (T, γ, τ) , where (T, γ) is a directed tree decomposition of κ and $\tau : \mathfrak{T} \rightarrow V(T)$ a bijective mapping such that the following two conditions are satisfied.

(DTD.1) For all nodes $t, u \in V(T)$ with $u \not\preceq t$ there is a minimum $(\tau^{-1}(u), \tau^{-1}(t))$ -separation Y such that $\gamma(u) \subseteq Y$.

(DTD.2) For all nodes $t \in V(T)$ except the root, there is a node $u \in V(T)$ such that $t \not\preceq u$ and $\gamma(t)$ is a leftmost minimum $(\tau^{-1}(t), \tau^{-1}(u))$ -separation.

Observe that (DTD.1) implies that for all nodes $t \in V(T)$ and children $u \in N_+(t)$ we have $\gamma(u) \not\subseteq \tau^{-1}(t)$. Furthermore, (DTD.2) implies that $\gamma(t) \in \tau^{-1}(t)$.

Recall that a κ -tangle \mathcal{T} is k -maximal, for some $k \geq 0$, if either $\text{ord}(\mathcal{T}) = k$ or $\text{ord}(\mathcal{T}) < k$ and \mathcal{T} is an (inclusionwise) maximal tangle. We denote the family of all k -maximal κ -tangles by $\mathfrak{T}_{\max}^{\leq k}$. Observe that for $k = \text{bw}(\kappa)$ the k -maximal κ -tangles are precisely the maximal κ -tangles.

Theorem 4.4 ([15]). *Let $\ell \geq 0$. Then there is a polynomial time algorithm that, given oracle access to a connectivity function κ and a κ -tangle $\mathcal{T}_{\text{root}} \in \mathfrak{T}_{\text{max}}^{\leq \ell}$ (via a membership oracle or its index in a comprehensive tangle data structure for κ), computes a canonical directed tree decomposition (T, γ, τ) for the set $\mathfrak{T}_{\text{max}}^{\leq \ell}$ such that $\tau^{-1}(r) = \mathcal{T}_{\text{root}}$ for the root r of T .*

Here canonical means that if $\kappa' : 2^{A'} \rightarrow \mathbb{N}$ is another connectivity function and $\mathcal{T}'_{\text{root}}$ an ℓ -maximal κ' -tangle, and (T', γ', τ') is the decomposition computed by our algorithm on input $(\kappa', \mathcal{T}'_{\text{root}})$, then for every isomorphism f from $(\kappa, \mathcal{T}_{\text{root}})$ to $(\kappa', \mathcal{T}'_{\text{root}})$, that is, bijective mapping $f : A \rightarrow A'$ with $\kappa(X) = \kappa'(f(X))$ and $X \in \mathcal{T}_{\text{root}} \iff f(X) \in \mathcal{T}'_{\text{root}}$ for all $X \subseteq A$, there is an isomorphism g from T to T' such that $f(\gamma(t)) = \gamma'(g(t))$ for all $t \in V(T)$ and $X \in \tau^{-1}(t) \iff f(X) \in (\tau')^{-1}(g(t))$ for all $X \subseteq A, t \in V(T)$.

5 Partitioning with Respect to a Maximal Tangle

Let G be a graph of rank width at most k . In this and the following two sections, we describe our construction of a canonical treelike decomposition of ρ_G of width at most $a(k)$ (for some function a). Since large parts of the construction go through for arbitrary connectivity functions, we find it convenient to let $\kappa := \rho_G$ and $A := V(G)$.

We start from a directed tree decomposition (T, γ, τ) for $\mathfrak{T}_{\text{max}}^{\leq k}$. The idea is to decompose the “pieces” of this decomposition, corresponding to the nodes of T , further into decompositions of bounded width and then merge all these bounded-width decompositions into one big decomposition. The largest part of the construction, resulting in Lemma 7.1, deals with a single node of T .

So we fix a node $t \in V(T)$. We let $\mathcal{T}_0 := \tau^{-1}(t)$ be the maximal tangle associated with t and $k_0 := \text{ord}(\mathcal{T}_0)$. Let $B := \beta(t)$ and $C_0 := \gamma(t)$. Assuming that the children of t in T are u_1, \dots, u_m , we let $C_i := \gamma(u_i)$ for $i \in [m]$. Observe that the sets B, C_0, \dots, C_m form a partition of A (the set C_0 may be empty). Now we contract the sets C_0, \dots, C_m . We shall construct a bounded width decomposition of the resulting connectivity function $\kappa \downarrow$ on the contracted set $A \downarrow$.

We construct the decomposition recursively. At any time, we have a set $X \subseteq A \downarrow$ that still needs to be decomposed, and we will show how to partition X in a canonical way, at any time keeping control of the width of the resulting decomposition.

We initialise the construction by taking a triple cover Q of the tangle \mathcal{T}_0 of size $|Q| \leq \theta(3k_0 - 2)$. We let Q^\vee be the “projection” of Q into $A \downarrow$ (precise definitions follow). The set $Q^\vee \cup \{c_0\}$ will be the bag at the root of our decomposition, and the first set X to be decomposed further is $A \downarrow \setminus (Q^\vee \cup \{c_0\})$.

Now suppose we are in some decomposition step where we need to decompose $X \subseteq A \downarrow \setminus (Q^\vee \cup \{c_0\})$. Depending on $\kappa \downarrow(X)$, we do this in two completely different ways. In this section (Section 5), we consider the case $\kappa \downarrow(X) < (3k + 2) \cdot k$, and in Section 6 we shall consider the case $\kappa \downarrow(X) \geq (3k + 2) \cdot k$.

5.1 Assumptions

Before we start the technical construction, we step back and collect the assumptions we make in a slightly more abstract setting, which we fix for the rest of the section.

Assumptions 5.1. (1) $\kappa : 2^A \rightarrow \mathbb{N}$ is a connectivity function on a set A .

(2) $k := \text{bw}(\kappa) \geq 1$.

(3) $C_0, \dots, C_m \subseteq A$ are mutually disjoint sets with $\kappa(C_i) < k$, and C_1, \dots, C_m are nonempty.

(4) For all $i \in [m]$ there are tangles $\mathcal{T}_i, \mathcal{T}'_i$ such that C_i is a leftmost minimum $(\mathcal{T}_i, \mathcal{T}'_i)$ -separation. Furthermore, if $C_0 \neq \emptyset$ then there are tangles $\mathcal{T}_0, \mathcal{T}'_0$ such that C_0 is a leftmost minimum $(\mathcal{T}_0, \mathcal{T}'_0)$ -separation.

$$(5) \ B := A \setminus (C_0 \cup \dots \cup C_m).$$

$$(6) \ A \downarrow := A \downarrow_{C_0, \dots, C_m}, \text{ and } c_i \text{ is the element of } A \downarrow \text{ corresponding to the contracted set } C_i, \text{ for } i = 0, \dots, m.$$

$$(7) \ \kappa \downarrow := \kappa \downarrow_{C_0, \dots, C_m}.$$

The assumption $\text{bw}(\kappa) \geq 1$ is without loss of generality, because if $\text{bw}(\kappa) = 0$ then $\kappa(\{x\}) = 0$ for all $x \in A$ and thus $\kappa(X) = 0$ for all $X \subseteq A$.

Assumptions 5.2. (1) There is a maximal κ -tangle \mathcal{T}_0 such that $\overline{C}_0 \in \mathcal{T}_0$ and $C_i \notin \mathcal{T}_0$ for $i = 1, \dots, m$.³

(2) For every κ -tangle $\mathcal{T} \perp \mathcal{T}_0$, there is an $i \in \{0, \dots, m\}$ and a set $Y \subseteq C_i$ such that $Y \in \mathcal{T}$.

(3) $k_0 := \text{ord}(\mathcal{T}_0)$. (Note that $k_0 \leq k$.)

Observe that

$$\mathcal{T}_0 \downarrow := \mathcal{T}_0 \downarrow_{C_0, \dots, C_m}.$$

is a $\kappa \downarrow$ -tangle, because $C_i \notin \mathcal{T}_0$ for $i = 0, \dots, m$ by Assumption 5.2(1).

Assumptions 5.3. (1) $Q \subseteq A$ is a triple cover of the tangle \mathcal{T}_0 of size $|Q| \leq \theta(3k_0 - 2)$.

(2) $Q^\vee := (B \cap Q) \cup \{c_i \mid 0 \leq i \leq m, C_i \cap Q \neq \emptyset\}$.

Observe that Q^\vee is a triple cover for the $\kappa \downarrow$ -tangle $\mathcal{T}_0 \downarrow$.

All algorithms we devise in this section will get κ and C_0, \dots, C_m and Q as input, and possibly other objects. We assume that we have constructed a comprehensive tangle data structure for κ and have determined the index of \mathcal{T}_0 in this data structure. Thus our algorithms also have access to \mathcal{T}_0 .

Whenever we refer to a construction in this section as being *canonical*, what we mean is that it is canonical given κ and C_0, \dots, C_m and Q . Note that \mathcal{T}_0 is canonical given κ and C_0, \dots, C_m , because \mathcal{T}_0 is the unique maximal κ -tangle with $\overline{C}_0 \in \mathcal{T}_0$ and $C_i \notin \mathcal{T}_0$ for $i = 1, \dots, m$. Thus we may depend on \mathcal{T}_0 in canonical constructions.

Our goal is to prove the following lemma.

Lemma 5.4. *For every $k_1 \in \mathbb{N}$ there are $a_1 = a_1(k, k_1)$, $b_1 = b_1(k, k_1)$, and $f_1 = f_1(k, k_1) > 0$ such that for every $X \subseteq A \downarrow \setminus (Q^\vee \cup \{c_0\})$ of order $\kappa \downarrow(X) = k_1$ and size $|X| \geq 2$, one of the following two conditions is satisfied.*

(i) *There is a canonical partition of X into $b \leq b_1$ sets X_1, \dots, X_b such that $\kappa \downarrow(X_i) \leq a_1$ and $|X_i| \leq (1 - 1/f_1)|X|$ for all $i \in [b]$.*

(ii) *There is a canonical partition of X into sets X_0, X_1, \dots, X_n such that*

- a. $\kappa \downarrow(X_0) \leq k_1$,
- b. $\kappa \downarrow(\bigcup_{i \in I} X_i) \leq 2k_1$ for every set $I \subseteq [n]$,
- c. $|X_i| \leq (1 - 1/f_1)|X|$ for every $i \in [n]$.

Furthermore, given X (in addition to $\kappa \downarrow$ and C_0, \dots, C_m and Q), the partition in (i) or (ii) can be computed in polynomial time (for fixed k, k_1).

³If $C_0 \neq \emptyset$, then the tangles \mathcal{T}_0 in Assumption 5.1(4) and Assumption 5.2(1) are the same, but this is irrelevant.

The lemma will be proved in Section 5.4.

For the rest of Section 5, we fix a set $X \subseteq A \downarrow \setminus (Q^\vee \cup \{c_0\})$. Let

$$\begin{aligned} k_1 &:= \kappa \downarrow(X), \\ k_2 &:= k_0 + k_1. \end{aligned}$$

We assume that

$$|X| \geq 6k_2. \quad (5.A)$$

Note that this implies $|X| \geq 6$ by our assumption that $k_0 \geq 1$.

5.2 Existence of a Balanced Separations

We call a set $Z \subseteq X$ a *balanced X -separation* if $\kappa \downarrow(Z) \leq k_1 = \kappa \downarrow(X)$ and

$$\frac{1}{3}|X| - k_2 + \kappa \downarrow(Z) \leq |Z| \leq \frac{2}{3}|X| + k_2 - \kappa \downarrow(Z)$$

Note that this notion does not only depend on X , but through k_2 also on k_0 , the order of the tangle \mathcal{T}_0 .

Lemma 5.5. *Suppose that $X \subseteq \overline{Y}$ for some Y in $\mathcal{T}_0 \downarrow$. Then there is a balanced X -separation.*

Proof. By the assumption of the lemma, there exists an $(X, \mathcal{T}_0 \downarrow)$ separation Y . Let X' be the leftmost minimum $(X, \mathcal{T}_0 \downarrow)$ -separation. Then $X \subseteq X' \subseteq A \downarrow$ and $\kappa \downarrow(X') \leq \kappa \downarrow(X) = k_1$ and $\overline{X'} \in \mathcal{T}_0 \downarrow$ and thus $\kappa \downarrow(X') < \text{ord}(\mathcal{T}_0 \downarrow) = k_0$.

Claim 1. There is a $Y \in \mathcal{T}_0 \downarrow$ such that $\kappa \downarrow(Y) \leq \kappa \downarrow(X')$ and

$$\frac{1}{3}|X| - k_0 + \kappa \downarrow(Y) \leq |X \cap Y| \leq \frac{2}{3}|X| + k_0 - \kappa \downarrow(Y).$$

Proof. We define a weight function $\varphi : A \rightarrow \mathbb{R}$ as follows:

$$\varphi(x) := \begin{cases} \frac{1}{|C_i|} & \text{if } x \in C_i \text{ for some } i \in [m] \text{ such that } c_i \in X, \\ 1 & \text{if } x \in X \cap B, \\ 0 & \text{otherwise.} \end{cases}$$

For a set $Y \subseteq A$ we let $\varphi(Y) := \sum_{y \in Y} \varphi(y)$. Note that for every $Y \subseteq A$ we have $\varphi(Y) + \varphi(\overline{Y}) = |X|$. Furthermore, for all $Y^\vee \subseteq A \downarrow$ we have $\varphi(Y^\vee \uparrow) = |X \cap Y^\vee|$.

Suppose first that there is no $Y \subseteq A$ such that $\kappa(Y) \leq \kappa \downarrow(X')$ and $\frac{1}{3}|X| \leq \varphi(Y) \leq \frac{2}{3}|X|$. Let

$$\mathcal{T} := \{Y \subseteq A \mid \kappa(Y) \leq \kappa \downarrow(X') \text{ and } \varphi(Y) > \frac{2}{3}|X|\}.$$

Then \mathcal{T} is a κ -tangle of order $\kappa \downarrow(X') + 1 \leq k_0$. To verify (T.3), let $x \in A$. Then $\varphi(\{x\}) \leq 1 \leq (2/3)|X|$, because $|X| \geq 2$.

Observe that $X' \uparrow \in \mathcal{T}$, because $\varphi(X' \uparrow) = |X \cap X'| = |X|$. Thus $\mathcal{T} \perp \mathcal{T}_0$.

For every $i \in \{0, \dots, m\}$ and every $Y \subseteq C_i$ we have $\varphi(Y) \leq \varphi(C_i) \leq 1$ and thus $Y \notin \mathcal{T}$. This contradicts Assumption 5.2(2).

We choose ℓ minimum such that there exists a $Y \subseteq A$ with $\kappa(Y) = \ell$ and

$$\frac{1}{3}|X| - k_0 + \ell \leq \varphi(Y) \leq \frac{2}{3}|X| + k_0 - \ell. \quad (5.B)$$

Then $\ell \leq \kappa \downarrow(X') \leq k_0$, because we have just proved that there is a set Y with $\kappa(Y) \leq \kappa \downarrow(X')$ and

$$\frac{1}{3}|X| - k_0 + \kappa(Y) \leq \frac{1}{3}|X| \leq \varphi(Y) \leq \frac{2}{3}|X| \leq \frac{2}{3}|X| + k_0 - \kappa(Y).$$

Without loss of generality we may assume that either $Y \cap C_0 = \emptyset$ or $C_0 \subseteq Y$. To see this, suppose that neither $Y \cap C_0 = \emptyset$ nor $C_0 \subseteq Y$, or equivalently, neither $Y \subseteq \overline{C_0}$ nor $\overline{Y} \subseteq \overline{C_0}$. By Assumption 5.1(4), $\overline{C_0}$ is a minimum $(\mathcal{T}_0, \mathcal{T}'_0)$ -separation. Thus by Lemma 2.1(1) (applied to $X = \overline{C_0}$), either $\kappa(Y \cap \overline{C_0}) \leq \kappa(Y)$ or $\kappa(\overline{Y} \cap \overline{C_0}) \leq \kappa(Y)$. As $c_0 \notin X$, we have $\varphi(Y \cap \overline{C_0}) = \varphi(Y \cup C_0) = \varphi(Y)$. Thus if $\ell' := \kappa(Y \cap \overline{C_0}) \leq \kappa(Y)$, then $Y' := Y \cap \overline{C_0}$ satisfies (5.B) with Y', ℓ' instead of Y, ℓ , and if $\ell' := \kappa(\overline{Y} \cap \overline{C_0}) \leq \kappa(Y)$, then $Y' := \overline{Y} \cap \overline{C_0}$ satisfies (5.B) with Y', ℓ' instead of Y, ℓ . In both cases, we have $Y' \cap C_0 = \emptyset$. This justifies the assumption that either $Y \cap C_0 = \emptyset$ or $C_0 \subseteq Y$.

Suppose for contradiction that there is an $i \in [m]$ such that neither $Y \cap C_i = \emptyset$ nor $C_i \subseteq Y$, or equivalently, neither $Y \subseteq \overline{C_i}$ nor $\overline{Y} \subseteq \overline{C_i}$. We argue similarly as for C_0 , but with Lemma 2.1(2). By Assumption 5.2, $\overline{C_i}$ is a rightmost minimum $(\mathcal{T}'_i, \mathcal{T}_i)$ -separation. Thus by Lemma 2.1(2) (applied to $X = \overline{C_i}$), either $\kappa(Y \cap \overline{C_i}) < \kappa(Y)$ or $\kappa(\overline{Y} \cap \overline{C_i}) < \kappa(Y)$. Without loss of generality, we assume that $\kappa(Y \cap \overline{C_i}) < \kappa(Y)$. The case $\kappa(\overline{Y} \cap \overline{C_i}) < \kappa(Y)$ is symmetric, because (5.B) is symmetric in Y, \overline{Y} . Let $Y' := Y \cap \overline{C_i} = Y \setminus C_i$ and $\ell' := \kappa(Y') < \ell$. We have

$$\varphi(Y) \geq \varphi(Y') = \varphi(Y) - \varphi(C_i) \geq \varphi(Y) - 1.$$

Thus by (5.B),

$$\begin{aligned} \frac{1}{3}|X| - k_0 + \ell' &\leq \frac{1}{3}|X| - k_0 + \ell - 1 \leq \varphi(Y) - 1 \\ &\leq \varphi(Y') \leq \varphi(Y) \leq \frac{2}{3}|X| + \kappa \downarrow(X') - \ell \\ &\leq \frac{2}{3}|X| + \kappa \downarrow(X') - \ell'. \end{aligned}$$

Thus $Y', \ell' < \ell$ satisfy (5.B). This contradicts the minimality of ℓ .

We have proved that for all $i \in [m]$, either $Y \cap C_i = \emptyset$ or $C_i \subseteq Y$. This implies that there is a $Y^\vee \subseteq A \downarrow$ such that $Y = Y^\vee \uparrow$. As $\varphi(Y) = |X \cap Y^\vee|$ and $\kappa \downarrow(Y^\vee) = \kappa(Y) = \ell$, by (5.B) we have

$$\frac{1}{3}|X| - k_0 + \kappa \downarrow(Y^\vee) \leq |X \cap Y^\vee| \leq \frac{2}{3}|X| + k_0 - \kappa \downarrow(Y^\vee). \quad (5.C)$$

As $\kappa \downarrow(Y^\vee) = \ell \leq \kappa \downarrow(X') < \text{ord}(\mathcal{T}_0 \downarrow)$, either $Y^\vee \in \mathcal{T}_0 \downarrow$ or $\overline{Y}^\vee \in \mathcal{T}_0 \downarrow$. Since (5.C) is symmetric in Y^\vee and \overline{Y}^\vee , we may assume that $Y^\vee \in \mathcal{T}_0 \downarrow$. \lrcorner

We choose $Y \in \mathcal{T}_0 \downarrow$ according to the claim. Let

$$Z := X \cap \overline{Y}.$$

Suppose for contradiction that $\kappa \downarrow(Z) > \kappa \downarrow(X)$. Then by submodularity, $\kappa \downarrow(X \cup \overline{Y}) < \kappa \downarrow(Y)$. As

$$\overline{X'} \cap Y \cap (X \cup \overline{Y}) \subseteq \overline{X'} \cap Y \cap (X' \cup \overline{Y}) = \emptyset,$$

we have $X \cup \overline{Y} \notin \mathcal{T}_0 \downarrow$ and thus $\overline{X \cup \overline{Y}} \in \mathcal{T}_0 \downarrow$, which means that $X \cup \overline{Y}$ is an $(X, \mathcal{T}_0 \downarrow)$ -separation. However, as $\kappa \downarrow(X \cup \overline{Y}) < \kappa \downarrow(Y) \leq \kappa \downarrow(X')$, this contradicts X' being a minimum $(X, \mathcal{T}_0 \downarrow)$ -separation.

Hence $\kappa \downarrow(Z) \leq \kappa \downarrow(X) = k_1$, and we have

$$\frac{1}{3}|X| - k_2 + \kappa \downarrow(Z) \leq \frac{1}{3}|X| - k_0 \leq \frac{1}{3}|X| - k_0 + \kappa \downarrow(Y) \leq |X \cap Y| = |X \setminus Z| = |X| - |Z|,$$

which implies $|Z| \leq \frac{2}{3}|X| + k_2 - \kappa \downarrow(Z)$. Similarly, $|Z| \geq \frac{1}{3}|X| - k_2 + \kappa \downarrow(Z)$. \square

In the following lemma, we drop the assumption $X \subseteq \overline{Y}$ for some $Y \in \mathcal{T}_0 \downarrow$.

Lemma 5.6. *There is a balanced X -separation.*

Proof. If $X \subseteq \overline{Y}$ for some $Y \in \mathcal{T}_0 \downarrow$, then the assertion follows from Lemma 5.5. So suppose that $X \cap Y \neq \emptyset$ for all $Y \in \mathcal{T}_0 \downarrow$.

Claim 1. There is a $Y \subseteq A \downarrow$ such that $\kappa \downarrow(Y) \leq k_0$ and

$$\frac{1}{3}|X| - k_0 + \kappa \downarrow(Y) \leq |X \cap Y| \leq \frac{1}{3}|X| - k_0 + \kappa \downarrow(Y).$$

Proof. As in the proof of the previous lemma, we define a weight function $\varphi : A \rightarrow \mathbb{R}$ by

$$\varphi(x) := \begin{cases} \frac{1}{|C_i|} & \text{if } x \in C_i \text{ for some } i \in [m] \text{ such that } c_i \in X, \\ 1 & \text{if } x \in X \cap B, \\ 0 & \text{otherwise.} \end{cases}$$

For a set $Y \subseteq A$ we let $\varphi(Y) := \sum_{y \in Y} \varphi(y)$.

Suppose first that there is no $Y \subseteq A$ such that $\kappa(Y) \leq k_0$ and $\frac{1}{3}|X| \leq \varphi(Y) \leq \frac{2}{3}|X|$. Let

$$\mathcal{T} := \{Y \subseteq A \mid \kappa(Y) \leq k_0 \text{ and } \varphi(Y) > \frac{2}{3}|X|\}.$$

Then \mathcal{T} is a κ -tangle of order $k_0 + 1$. We have $\mathcal{T} \perp \mathcal{T}_0$, because \mathcal{T}_0 is maximal and $\text{ord}(\mathcal{T}) > \text{ord}(\mathcal{T}_0)$. For every $i \in \{0, \dots, m\}$ and every $Y \subseteq C_i$ we have $\varphi(Y) \leq \varphi(C_i) \leq 1$ and thus $Y \notin \mathcal{T}$. This contradicts Assumption 5.2(2).

We choose ℓ minimum such that there exists a $Y \subseteq A$ with $\kappa(Y) = \ell$ and

$$\frac{1}{3}|X| - k_0 + \ell \leq \varphi(Y) \leq \frac{2}{3}|X| + k_0 - \ell. \quad (5.D)$$

Arguing as in the proof of Claim 1 in the proof of Lemma 5.5, we prove that $\ell \leq k_0$ and that we may assume $C_i \cap Y = \emptyset$ or $C_i \subseteq Y$ for all $i \in \{0, \dots, m\}$.

Then there is a $Y^\vee \subseteq A \downarrow$ such that $Y = Y^\vee \uparrow$. As $\varphi(Y) = |X \cap Y^\vee|$, the claim follows. \square

We choose $Y \subseteq A \downarrow$ according to the claim. Suppose for contradiction that $\kappa \downarrow(X \cap Y) > \kappa \downarrow(X)$ and $\kappa \downarrow(X \cap \overline{Y}) > \kappa \downarrow(X)$. Then by submodularity, $\kappa \downarrow(X \cup Y) < \kappa \downarrow(Y) \leq k_0$ and $\kappa \downarrow(X \cup \overline{Y}) < k_0$. As X has a nonempty intersection with every element of $\mathcal{T}_0 \downarrow$, we have $X \cup Y, X \cup \overline{Y} \in \mathcal{T}_0 \downarrow$. Thus $Q^\vee \cap (X \cup Y) \cap (X \cup \overline{Y}) \neq \emptyset$, because Q^\vee is a triple cover of $\mathcal{T}_0 \downarrow$ (a double cover would be sufficient here). This is a contradiction.

Thus either $\kappa \downarrow(X \cap Y) \leq \kappa \downarrow(X)$ or $\kappa \downarrow(X \cap \overline{Y}) \leq \kappa \downarrow(X)$. Without loss of generality we assume the former. Then it is easy to see that $Z := X \cap Y$ is X -balanced. \square

5.3 A Canonical Family of Separations

For $0 \leq \ell \leq k_1$, let

$$p(\ell) := 2^{-3^{k_2 - \ell}}.$$

Note that

$$\frac{1}{8} \geq p(k_1) \geq p(k_1 - 1) \geq \dots \geq p(0) \quad (5.E)$$

and

$$p(\ell - 1) = p(\ell)^3 \quad (5.F)$$

for all $\ell \geq 1$.

Let us call a set $Z \subseteq X$ of order $\ell := \kappa \downarrow(Z)$ *good* (or a *good separation*) if

$$p(\ell) \cdot |X| \leq |Z| < |X|. \quad (5.G)$$

Recall that $\frac{1}{6}|X| \geq k_2$ by (5.A). Thus for $0 \leq \ell \leq k_1$ we have

$$\frac{1}{3}|X| - k_2 + \ell \geq \frac{1}{3}|X| - k_2 \geq \frac{1}{6}|X| \geq \max \{1, p(\ell) \cdot |X|\}$$

It follows that every balanced X -separation is good. Hence by Lemma 5.6, there is a good separation Z of order $\kappa\downarrow(Z) \leq k_1$.

Let ℓ be minimum such that there is a good separation Z of order $\kappa\downarrow(Z) = \ell$.

Let \mathcal{Z} be the set of all $Z \subseteq X$ such that

- (i) Z is good;
- (ii) $\kappa\downarrow(Z) = \ell$;
- (iii) $|Z|$ is maximum subject to (i) and (ii).

Observe that $|Z| = |Z'| \geq p(\ell) \cdot |X|$ for all $Z, Z' \in \mathcal{Z}$. Let

$$\mathcal{Y} := \{Y \subseteq A \downarrow \mid \overline{Y} \in \mathcal{Z}\}.$$

Note that $\overline{X} \subset Y$ and $\kappa\downarrow(Y) = \ell$ and $|Y| = |Y'|$ for all $Y, Y' \in \mathcal{Y}$. Let us call two sets $Y, Y' \in \mathcal{Y}$ X -disjoint if $Y \cap Y' \subseteq \overline{X}$. Observe that for X -disjoint sets $Y, Y' \in \mathcal{Y}$ we have $Y \cap Y' = \overline{X}$.

Our next goal is to prove the following lemma.

Lemma 5.7. *There is a $b_2 = b_2(k_0, k_1)$ such that if $|\mathcal{Y}| > b_2$ then the elements of \mathcal{Y} are mutually X -disjoint.*

The idea of the proof is as follows. Assume that there are $Y, Y' \in \mathcal{Y}$ that are not X -disjoint. Then there are $Z, Z' \in \mathcal{Z}$ whose union is a proper subset of X . The choice of the function p and a submodularity argument guarantee that these sets Z, Z' have a small intersection. Thus $|Y \setminus Y'| = |Z' \setminus Z|$ is relatively large (close to $p(\ell)|X|$, i.e., a constant fraction of $|X|$) and thus $|Y \setminus \overline{X}| \geq |Y \setminus Y'|$ is relatively large. As all elements of \mathcal{Y} have the same size, this holds for all $Y \in \mathcal{Y}$. Now we apply Ramsey's Theorem and find that if \mathcal{Y} is very large either (i) there is a large family $Z_1, \dots, Z_n \in \mathcal{Z}$ such that all pairwise unions $Z_i \cup Z_j$ are proper subsets of X , or (ii) there is a large family $Z_1, \dots, Z_n \in \mathcal{Z}$ such that all pairwise unions $Z_i \cup Z_j$ are equal to X . In case (i), we argue that the Z_i are relatively large, but have a small intersection, and thus for large n their union becomes larger than $|X|$, which is impossible. In case (ii) we argue that the $Y_i := \overline{Z_i}$ are mutually X -disjoint, and as the sets $Y_i \setminus \overline{X}$ are relatively large, for large n their union becomes larger than $|X|$. Again, this is impossible. Thus the size of \mathcal{Y} must be bounded.

The actual proof requires some preparation.

Lemma 5.8. *Let $Z, Z' \in \mathcal{Z}$ with $Z \neq Z'$ and $Z \cup Z' \neq X$. Then*

$$|Z \cap Z'| < p(\ell)^3 \cdot |X|$$

Proof. Let $\ell_{\cup} := \kappa\downarrow(Z \cup Z')$. Suppose for contradiction that $\ell_{\cup} \leq \ell$. Then

$$p(\ell_{\cup}) \cdot |X| \leq p(\ell) \cdot |X| \leq |Z| \leq |Z \cup Z'| < |X|,$$

where the last inequality holds, because $Z \cup Z' \neq X$. By the choice of ℓ we thus have $\ell_{\cup} = \ell$. Hence $Z \cup Z'$ satisfies (i) and (ii). However, $|Z_{\cup}| > |Z|$, because $Z \neq Z'$ and $|Z| = |Z'|$. This is a contradiction, which proves $\kappa\downarrow(Z \cup Z') = \ell_{\cup} > \ell$.

By submodularity, $\ell_{\cap} := \kappa\downarrow(Z \cap Z') < \ell$. Thus by the choice of ℓ we have

$$|Z \cap Z'| < p(\ell_{\cap}) \cdot |X| \leq p(\ell)^3 \cdot |X|. \quad \square$$

Lemma 5.9. *Let $Z_1, \dots, Z_m \in \mathcal{Z}$ such that $Z_i \cup Z_j \neq X$ for all distinct $i, j \in [m]$. Then*

$$m < \frac{2}{p(\ell)}$$

Proof. As $|Z_i| \geq p(\ell) \cdot |X|$ and $|Z_i \cap Z_j| \leq p(\ell)^3 \cdot |X|$, we have

$$\left| Z_i \setminus \bigcup_{j=1}^{i-1} Z_j \right| \geq (p(\ell) - (i-1) \cdot p(\ell)^3) \cdot |X|.$$

This implies, for all $i \leq m$

$$|X| \geq \left| \bigcup_{j=1}^i Z_j \right| \geq \left(i \cdot p(\ell) - \sum_{j=1}^{i-1} j \cdot p(\ell)^3 \right) \cdot |X|.$$

Thus

$$\begin{aligned} \left(i \cdot p(\ell) - \frac{i \cdot (i-1)}{2} \cdot p(\ell)^3 \right) &= \left(i \cdot p(\ell) - \sum_{j=0}^{i-1} j \cdot p(\ell)^3 \right) \leq 1 \\ \iff \frac{p(\ell)^3}{2} \cdot i^2 - \left(p(\ell) + \frac{p(\ell)^3}{2} \right) \cdot i + 1 &\geq 0 \end{aligned}$$

It is easy to see that this last inequality is violated for $i = \frac{2}{p(\ell)}$, which is an integer because $\frac{1}{p(\ell)}$ is a power of 2. Thus $m < \frac{2}{p(\ell)}$. \square

Lemma 5.10. *For all $n \geq 1$ there is an $m = m(\ell, n)$ such that if $|\mathcal{Z}| > m$ then there are $Z_1, \dots, Z_n \in \mathcal{Z}$ such that $Z_i \cup Z_j = X$ for all distinct $i, j \in [n]$.*

Proof. If \mathcal{Z} is sufficiently large, then by Ramsey's Theorem one of the following two assertions holds.

- (i) There are $Z_1, \dots, Z_{2/p(\ell)} \in \mathcal{Z}$ such that $Z_i \cup Z_j \neq X$ for all distinct $i, j \in [2/p(\ell)]$.
- (ii) There are $Z_1, \dots, Z_n \in \mathcal{Z}$ such that $Z_i \cup Z_j = X$ for all distinct $i, j \in [n]$.

Lemma 5.9 rules out (i). Thus (ii) holds. \square

Let us now turn to the sets in \mathcal{Y} . Recall that $Y, Y' \in \mathcal{Y}$ are X -disjoint if $Y \cap Y' = \overline{X}$.

Lemma 5.11. *If there are distinct sets in \mathcal{Y} that are not X -disjoint, then for all $Y \in \mathcal{Y}$,*

$$|Y \setminus \overline{X}| \geq (p(\ell) - p(\ell)^3) \cdot |X|.$$

Proof. Let $Y_1, Y_2 \in \mathcal{Y}$ such that $Y_1 \neq Y_2$ and $Y_1 \cap Y_2 \neq \overline{X}$, and let $Z_1 := \overline{Y_1}$ and $Z_2 := \overline{Y_2}$. Then $Z_1 \cup Z_2 \neq X$. Thus

$$|Y_1 \setminus \overline{X}| \geq |Y_1 \setminus Y_2| = |Z_2 \setminus Z_1| = |Z_2| - |Z_1 \cap Z_2| \geq (p(\ell) - p(\ell)^3) \cdot |X|,$$

where the last inequality follows from Lemma 5.8. For all $Y \in \mathcal{Y}$ we have $|Y| = |Y_1|$ and thus $|Y \setminus \overline{X}| = |Y_1 \setminus \overline{X}|$. \square

Proof of Lemma 5.7. We let

$$n(\ell) := \left\lfloor \frac{1}{(p(\ell) - p(\ell)^3)} \right\rfloor + 1,$$

and choose $m = m(\ell, n(\ell))$ according to Lemma 5.10. Suppose that $|\mathcal{Z}| = |\mathcal{Y}| > m$. Then there are $Z_1, \dots, Z_{n(\ell)}$ such that $Z_i \cup Z_j = X$ for all distinct $i, j \in [n(\ell)]$. For all $i \in [n(\ell)]$, let $Y_i = \overline{Z_i}$. Then $Y_1, \dots, Y_{n(\ell)}$ are mutually X -disjoint. Thus

$$\sum_{i=1}^{n(\ell)} |Y_i \setminus \overline{X}| \leq |X|.$$

By the choice of $n(\ell)$ and Lemma 5.11, it follows that all sets in \mathcal{Y} are mutually X -disjoint.

To complete the proof, we let

$$b_2(k_0, k_1) := \max_{0 \leq \ell \leq k_1} m(\ell, n(\ell)).$$

□

Lemma 5.12. *Suppose that the elements of \mathcal{Y} are mutually X -disjoint. Then for all $\mathcal{Y}_0 \subseteq \mathcal{Y}$,*

$$\kappa \downarrow \left(\bigcup_{Y \in \mathcal{Y}_0} Y \right) \leq \ell,$$

with equality for all $\mathcal{Y}_0 \subset \mathcal{Y}$.

Proof. We prove by induction on $i \leq |\mathcal{Y}|$ that for all $Y_1, \dots, Y_i \in \mathcal{Y}$,

$$\kappa \downarrow (Y_1 \cup \dots \cup Y_i) \leq \ell, \tag{5.H}$$

with equality if $i < |\mathcal{Y}|$.

The base step $i = 1$ is trivial. For $i = 2$, let $Y_1, Y_2 \in \mathcal{Y}$. Then $Y_1 \cap Y_2 = \overline{X}$ and thus $\kappa \downarrow (Y_1 \cap Y_2) = \kappa \downarrow (X) = k_1 \geq \ell$. By submodularity, $\kappa \downarrow (Y_1 \cup Y_2) \leq \ell$.

Now let $2 \leq i < |\mathcal{Y}|$, and suppose that

$$\kappa \downarrow (Y_1 \cup \dots \cup Y_{i-1}) = \ell \tag{5.I}$$

for all $Y_1, \dots, Y_{i-1} \in \mathcal{Y}$ and (5.H) for all $Y_1, \dots, Y_i \in \mathcal{Y}$.

Let $Y_1, \dots, Y_{i+1} \in \mathcal{Y}$ be mutually distinct. By posimodularity

$$\kappa \downarrow (Y_1 \cup \dots \cup Y_i) + \kappa \downarrow (Y_i \cup Y_{i+1}) \geq \kappa \downarrow (Y_1 \cup \dots \cup Y_{i-1}) + \kappa \downarrow (Y_{i+1}).$$

As $\kappa \downarrow (Y_1 \cup \dots \cup Y_{i-1}) = \kappa \downarrow (Y_{i+1}) = \ell$ by (5.I) and $\kappa \downarrow (Y_i \cup Y_{i+1}) \leq \ell$ by (5.H), it follows that $\kappa \downarrow (Y_1 \cup \dots \cup Y_i) \geq \ell$, which combined with (5.H) implies equality.

Furthermore, by submodularity,

$$\kappa \downarrow (Y_1 \cup \dots \cup Y_i) + \kappa \downarrow (Y_i \cup Y_{i+1}) \geq \kappa \downarrow (Y_i) + \kappa \downarrow (Y_1 \cup \dots \cup Y_{i+1}).$$

As $\kappa \downarrow (Y_1 \cup \dots \cup Y_i) = \kappa \downarrow (Y_i \cup Y_{i+1}) = \kappa \downarrow (Y_i) = \ell$, this implies $\kappa \downarrow (Y_1 \cup \dots \cup Y_{i+1}) \leq \ell$. □

5.4 Proof of Lemma 5.4

We continue to use the notation of Section 5.3. Essentially, the lemmas proved there show how to use the family \mathcal{Y} to obtain the desired partition of X . The main question that remains to be solved is how to compute \mathcal{Y} .

Lemma 5.13. *There is a polynomial time algorithm that, given X and oracle access to $\kappa \downarrow$, computes \mathcal{Y} .*

Proof. Let \mathcal{Z}^* be the family of all $Z \subseteq A$ satisfying the following three conditions:

- (i) Z satisfies (5.G), that is, $p(\ell) \cdot |X| \leq |Z| < |X|$.
- (ii') $\kappa \downarrow_{\min}(Z, \overline{X}) = \ell$;
- (iii') there are a set $Z_0 \subseteq X$ of size $|Z_0| \leq \ell$ and an element $x \in X$ such that Z is a rightmost minimum $(Z_0, \overline{X} \cup \{x\})$ -separation.

Let $m := \max \{|Z| \mid Z \in \mathcal{Z}^*\}$.

Claim 1.

$$\mathcal{Z} = \{Z \in \mathcal{Z}^* \mid |Z| = m\}.$$

Proof. We first prove that $\mathcal{Z} \subseteq \mathcal{Z}^*$. Let $Z \in \mathcal{Z}$. Clause (i) in the definition of \mathcal{Z} is the same as clause (i) above.

If there was some Z' such that $Z \subseteq Z' \subseteq X$ and $\kappa \downarrow(Z') < \ell$, then $Z' \subset X$, because $\kappa \downarrow(X) = k_1 \geq \ell$, and Z' would also satisfy (5.G), because $|Z'| \geq |Z|$. Thus (ii) would be violated. This proves (ii').

To see that Z satisfies (iii'), let $Z_0 \subseteq Z$ be inclusionwise minimal such that $\kappa \downarrow_{\min}(Z_0, \overline{X}) = \ell$. Suppose for contradiction that $|Z_0| = n > \ell$, and let z_1, \dots, z_n be an enumeration of Z_0 . For every $i \in [n]$, let $Z^i = \{z_1, \dots, z_i\}$. Then $\kappa \downarrow_{\min}(Z^i, \overline{X}) \leq \kappa \downarrow_{\min}(Z^{i+1}, \overline{X})$ for all $i < \ell$, because $\kappa \downarrow_{\min}$ is monotone in the first argument. As $\kappa \downarrow_{\min}(Z_0, \overline{X}) = \ell$, there is an $i < \ell$ such that $\kappa \downarrow_{\min}(Z^i, \overline{X}) = \kappa \downarrow_{\min}(Z^{i+1}, \overline{X})$. By the submodularity of $\kappa \downarrow_{\min}$ in the first argument,

$$\kappa \downarrow_{\min}(Z_0 \setminus \{z_{i+1}\}, \overline{X}) + \kappa \downarrow_{\min}(Z^{i+1}, \overline{X}) \geq \kappa \downarrow_{\min}(Z^i, \overline{X}) + \kappa \downarrow_{\min}(Z_0, \overline{X}).$$

It follows that $\kappa \downarrow_{\min}(Z_0 \setminus \{z_{i+1}\}, \overline{X}) = \kappa \downarrow_{\min}(Z_0, \overline{X})$, contradicting the minimality of Z_0 . This proves that $|Z_0| \leq \ell$.

Let $x \in X \setminus Z$. Then $Z_0 \subseteq Z \subseteq X \setminus \{x\} \subseteq X$, and $\kappa \downarrow_{\min}(Z_0, \overline{X}) = \ell$ and $\kappa(Z) = \ell$ imply $\kappa \downarrow_{\min}(Z_0, \overline{X} \cup \{x\}) = \ell$. Thus Z is a minimum $(Z_0, \overline{X} \cup \{x\})$ -separation, and now clause (iii) in the definition of \mathcal{Z} (the maximality of $|Z|$) implies that Z is rightmost. This completes the proof of (iii') and thus of the inclusion $\mathcal{Z} \subseteq \mathcal{Z}^*$. The maximality of the elements of \mathcal{Z} (clause (iii) in the definition) then implies that

$$\mathcal{Z} \subseteq \{Z \in \mathcal{Z}^* \mid |Z| = m\}.$$

To prove the converse inclusion, let $Z \in \mathcal{Z}^*$ with $|Z| = m$. Then by (iii'), $Z \subset \overline{X}$. Clauses (i) and (ii') above imply clauses (i) and (ii) in the definition of \mathcal{Z} .

Suppose that there is some $Z' \subset X$ satisfying (i) and (ii) such that $|Z'| > |Z|$. Choose such a Z' of maximum size. Then $Z' \in \mathcal{Z}$, and thus $|Z'| = m = |Z|$. This is a contradiction. \lrcorner

It is easy to see that \mathcal{Z}^* can be computed in polynomial time, and this implies that \mathcal{Z} and thus \mathcal{Y} can be computed in polynomial time. \square

Proof of Lemma 5.4. Recall that $k_0 \leq k$. We let

$$f_1 := f_1(k, k_1) := \frac{1}{p(0) - p(0)^3} \tag{5.J}$$

and

$$a_1 := a_1(k, k_1) := \max\{k, 2k_1 \cdot b_2(k, k_1)\}, \tag{5.K}$$

$$b_1 := b_1(k, k_1) := \max\{6(k + k_1), 2^{b_2(k, k_1)}\}, \tag{5.L}$$

where $b_2 := b_2(k, k_1)$ is chosen according to Lemma 5.7.

If $|X| < 6(k + k_1)$, we simply partition X into 1-element sets. Note that $\kappa \downarrow(\{b\}) = \kappa(\{b\}) \leq k$ for all $b \in B$, because $\text{bw}(\kappa) \leq k \leq a_1$, and $\kappa(\{c_i\}) = \kappa(C_i) < k \leq a_1$ for $0 \leq i \leq m$ by Assumption 5.1(3). Thus (i) is satisfied.

In the following, we assume that $|X| \geq 6(k + k_1) \geq 6k_2$. This is the assumption needed for the previous results.

Case 1: There are distinct $Y, Y' \in \mathcal{Y}$ that are not X -disjoint.

Then $|\mathcal{Y}| \leq b_2(k, k_2)$ by Lemma 5.7 and

$$|Y \cap X| > (p(\ell) - p(\ell)^3)|X| \geq \frac{1}{f_1} \cdot |X| \quad (5.M)$$

by Lemma 5.11. Moreover, $\overline{Y} \in \mathcal{Z}$ for all $Y \in \mathcal{Y}$, which implies

$$|\overline{Y}| \geq p(\ell) \cdot |X| \geq \frac{1}{f_1} \cdot |X|. \quad (5.N)$$

Let Y_1^+, \dots, Y_n^+ be an enumeration of all sets $|Y \cap X|$ for $Y \in \mathcal{Y}$. Note that $n \leq b_2(k, k_1)$. For every $i \in [n]$, let $Y_i^- := X \setminus Y_i^+$. For every $i \in [n]$, we have $\kappa \downarrow(Y_i^-) = \ell$ and, by submodularity, $\kappa \downarrow(Y_i^+) \leq k_1 + \ell \leq 2k_1$.

Let X_1, \dots, X_b be a list of all nonempty sets of the form

$$\bigcap_{i=1}^n Y_i^{\sigma(i)},$$

for some function $\sigma : [n] \rightarrow \{+, -\}$. Then $b \leq 2^n \leq b_1$. Submodularity implies that

$$\kappa \downarrow(X_i) \leq 2k_1 n \leq a_1.$$

It follows from (5.M) and (5.N) that $|Y_i^\sigma| \leq (1 - 1/f_1)|X|$ for all $i \in [n]$ and $\sigma \in \{+, -\}$. Thus the partition X_1, \dots, X_b satisfies assertion (i) of the lemma.

Case 2: The elements of \mathcal{Y} are mutually X -disjoint.

We let

$$X_0 := \bigcap_{Y \in \mathcal{Y}} \overline{Y},$$

and we let X_1, \dots, X_n be an enumeration of the sets $Y \cap X$ for $Y \in \mathcal{Y}$. Note that the sets X_0, \dots, X_n form a partition of X . It follows from Lemma 5.12 that $\kappa \downarrow(X_0) \leq \ell \leq k$ and $\kappa \downarrow(\bigcup_{i \in I} X_i) \leq k + \ell \leq 2k$ for every set $I \subseteq [n]$. It follows from (5.G) that $|X_i| \leq (1 - 1/f_1)|X|$ for every $i \in [n]$.

Thus the partition X_0, \dots, X_n satisfies assertion (i) of the lemma.

It follows from Lemma 5.13 that in both cases the partition can be computed in polynomial time. \square

6 The Non-Well-Linked Case

6.1 Partitioning with Respect to an Independent Set

In this section, we make Assumptions 5.1 again (but not Assumptions 5.2 and 5.3).

Let $X \subseteq A \downarrow$ such that

$$k_1 := \kappa \downarrow(X) \geq (3k + 2) \cdot k. \quad (6.A)$$

We define a function $\lambda : 2^{\overline{X}} \rightarrow \mathbb{N}$ by letting

$$\lambda(Y) := \kappa \downarrow_{\min}(Y, X)$$

for all $Y \subseteq \overline{X}$. Then λ is submodular and monotone, and we have $\lambda(\emptyset) = 0$. Such a function is known as an *integer polymatroid*. It induces a matroid $\mathcal{M}(\lambda)$ on \overline{X} whose independent sets are all $Y \subseteq \overline{X}$ satisfying

$$|Z| \leq \lambda(Z) \text{ for all } Z \subseteq Y \quad (6.B)$$

(see [27], Proposition 12.1.2). The rank function r_λ of $\mathcal{M}(\lambda)$ is defined by

$$r_\lambda(Y) := \min \{ \lambda(Z) + |Y \setminus Z| \mid Z \subseteq Y \}.$$

(see [27], Proposition 12.1.7). Observe that for all $y \in \overline{X}$ we have $\lambda(\{y\}) \leq \kappa \downarrow(\{y\}) \leq k$. If $y \in B$ then this holds because $\kappa \downarrow(\{y\}) = \kappa(\{y\}) \leq \text{bw}(\kappa)$, and if $y = c_i$ it follows from Assumption 5.1(3). A straightforward induction based on the submodularity of λ then implies that $\lambda(Y) \leq k|Y|$ for all $Y \subseteq \overline{X}$. Thus for all $Z \subseteq \overline{X}$,

$$\lambda(Z) \geq \lambda(\overline{X}) - \lambda(\overline{X} \setminus Z) \geq k_1 - k \cdot |\overline{X} \setminus Z|,$$

which implies $\lambda(Z) + k|\overline{X} \setminus Z| \geq k_1$ and hence $\lambda(Z) + |\overline{X} \setminus Z| \geq k_1/k$. By the definition of r_λ , we get

$$r_\lambda(\overline{X}) \geq \frac{k_1}{k} \geq (3k+2).$$

Thus there is a set $Y' \subseteq \overline{X}$ of size $|Y'| = 3k+2$ that is an independent set of $\mathcal{M}(\lambda)$. As all subsets of an independent set are independent as well, there is an independent set $Y \subseteq \overline{X} \setminus \{c_0\}$ of size $|Y| = 3k+1$. We keep such a set Y fixed in the following.

Lemma 6.1. *Let $Z \subseteq A \downarrow$ such that $\kappa \downarrow(Z) < |Y \setminus Z|$.*

Then

$$\kappa \downarrow(X \cap Z) < \kappa \downarrow(X).$$

Proof. As Y is independent, we have $|Y \setminus Z| \leq \lambda(Y \setminus Z)$. As $Y \setminus Z \subseteq \overline{X} \setminus Z \subseteq \overline{X}$, we have $\lambda(Y \setminus Z) \leq \kappa \downarrow(\overline{X} \setminus Z)$. Thus $\kappa \downarrow(\overline{X} \setminus Z) - |Y \setminus Z| \geq 0$ and therefore

$$\begin{aligned} \kappa \downarrow(X \cap Z) &\leq \kappa \downarrow(X \cap Z) + \kappa \downarrow(\overline{X} \setminus Z) - |Y \setminus Z| \\ &= \kappa \downarrow(X \cap Z) + \kappa \downarrow(X \cup Z) - |Y \setminus Z| && \text{(symmetry)} \\ &\leq \kappa \downarrow(X) + \kappa \downarrow(Z) - |Y \setminus Z| && \text{(submodularity)} \\ &< \kappa \downarrow(X) + |Y \setminus Z| - |Y \setminus Z| && \text{(assumption of the lemma)} \\ &= \kappa \downarrow(X). \end{aligned}$$

□

Lemma 6.2. *There is a set $Z \subseteq A \downarrow$ such that $\kappa \downarrow(Z) \leq k$ and*

$$\kappa \downarrow(Z) < \min\{|Y \cap Z|, |Y \setminus Z|\}. \quad (6.C)$$

Furthermore, we can compute such a set Z in polynomial time (for fixed k).

Proof. We define a weight function $\varphi : A \rightarrow \mathbb{R}$ as follows:

$$\varphi(x) := \begin{cases} \frac{1}{|C_i|} & \text{if } x \in C_i \text{ for some } i \in \{0, \dots, m\} \text{ such that } c_i \in Y, \\ 1 & \text{if } x \in Y \cap B, \\ 0 & \text{otherwise.} \end{cases}$$

Claim 1. There is a $Z \subseteq A$ such that $\kappa(Z) \leq k$ and

$$\min\{\varphi(Y \uparrow \cap Z), \varphi(Y \uparrow \setminus Z)\} > k.$$

Proof. Suppose for contradiction that for all $Z \subseteq A$ such that $\kappa(Z) \leq k$, either $\varphi(Y \uparrow \cap Z) \leq k$ or $\varphi(Y \uparrow \setminus Z) \leq k$. Let

$$\mathcal{T} := \{Z \subseteq A \mid \kappa(Z) \leq k, \varphi(Y \uparrow \cap Z) \geq 2k+1\}.$$

Then \mathcal{T} is a κ -tangle of order $k + 1$. Indeed, it obviously satisfies (T.0). It satisfies (T.1), because

$$\varphi(Y \uparrow \cap Z) + \varphi(Y \uparrow \setminus Z) = \varphi(Y \uparrow) = |Y| = 3k + 1.$$

It satisfies (T.2), because $2k + 1 > (2/3)|Y|$, and it satisfies (T.3) because $\varphi(\{x\}) \leq 1 < 2k + 1$ for all $x \in A$.

However, as $\text{bw}(\kappa) = k$ no tangle of order $k + 1$ exists. \lrcorner

Let ℓ be minimum such that there is a $Z \subseteq A$ such that $\kappa(Z) \leq \ell$ and

$$\min\{\varphi(Y \uparrow \cap Z), \varphi(Y \uparrow \setminus Z)\} > \ell. \quad (6.D)$$

Let $Z \subseteq A$ such that $\kappa(Z) \leq \ell$ and (6.D).

Without loss of generality we may assume that either $Z \cap C_0 = \emptyset$ or $C_0 \subseteq Z$. To see this, suppose that neither $Z \cap C_0 = \emptyset$ nor $C_0 \subseteq Z$, or equivalently, neither $Z \subseteq \overline{C_0}$ nor $\overline{Z} \subseteq \overline{C_0}$. By Assumption 5.1(4), $\overline{C_0}$ is a minimum $(\mathcal{T}_0, \mathcal{T}'_0)$ -separation. Thus by Lemma 2.1(1) (applied to $X = \overline{C_0}$), either $\kappa(Z \cap \overline{C_0}) \leq \kappa(Z)$ or $\kappa(\overline{Z} \cap \overline{C_0}) \leq \kappa(Z)$. As $c_0 \notin Y$, we have $\varphi(Z \cap \overline{C_0}) = \varphi(Z \cup C_0) = \varphi(Z)$. Thus if $\ell' := \kappa(Z \cap \overline{C_0}) \leq \kappa(Z)$, then $Z' := Z \cap \overline{C_0}$ satisfies (6.D) with Z', ℓ' instead of Z, ℓ , and if $\ell' := \kappa(\overline{Z} \cap \overline{C_0}) \leq \kappa(Z)$, then $Z' := \overline{Z} \cap \overline{C_0}$ satisfies (6.D) with Z', ℓ' instead of Z, ℓ . In both cases, we have $Z' \cap C_0 = \emptyset$. This justifies the assumption that either $Z \cap C_0 = \emptyset$ or $C_0 \subseteq Z$.

Claim 2. For all $i \in [m]$, either $C_i \cap Z = \emptyset$ or $C_i \subseteq Z$.

Proof. Suppose for contradiction that neither $C_i \cap Z = \emptyset$ nor $C_i \subseteq Z$. Then neither $Z \subseteq \overline{C_i}$ nor $\overline{Z} \subseteq \overline{C_i}$. By Assumption 5.1(4) and Lemma 2.1, either $\kappa(Z \cap \overline{C_i}) < \kappa(Y)$ or $\kappa(\overline{Z} \cap \overline{C_i}) < \kappa(Y)$.

Without loss of generality we assume that $\ell' := \kappa(Z \cap \overline{C_i}) < \kappa(Y) = \ell$. Let $Z' := Z \cap \overline{C_i}$. Then

$$\varphi(Y \uparrow \cap Z') \geq \varphi(Y \uparrow \cap Z) - 1 \geq \ell - 1 \geq \ell'$$

and

$$\varphi(Y \uparrow \setminus Z') \geq \varphi(Y \uparrow \setminus Z) \geq \ell \geq \ell'.$$

This contradicts the minimality of ℓ . \lrcorner

It follows that there is a $Z^\vee \subseteq A \downarrow$ such that $Z = Z^\vee \uparrow$. Then

$$|Y \cap Z^\vee| = \varphi(Y \uparrow \cap Z) > \ell = \kappa(Z) = \kappa \downarrow(Z^\vee)$$

and, similarly, $|Y \setminus Z^\vee| > \kappa \downarrow(Z^\vee)$.

We can compute a set Z satisfying (6.C) in polynomial time as follows: for every $Z_0 \subseteq Y$ we compute a leftmost minimum $(Z_0, Y \setminus Z_0)$ separation Z until we find one with $\kappa \downarrow(Z) < |Z_0| = |Y \cap Z|$ and $\kappa \downarrow(Z) < |Y \setminus Z_0| = |Y \setminus Z|$. \square

Lemma 6.3. *Let $Z \subseteq A \downarrow$ such that $\kappa \downarrow(Z) < \min\{|Y \cap Z|, |Y \setminus X|\}$. Then $X \cap Z, X \setminus Z$ is a partition of X into two nonempty sets with $\kappa \downarrow(X \cap Z), \kappa \downarrow(X \setminus Z) < \kappa \downarrow(X)$.*

Proof. As Y is independent in $\mathcal{M}(\lambda)$, for each Z with $\kappa \downarrow(Z) < \min\{|Y \cap Z|, |Y \setminus X|\}$ we have $X \cap Z \neq \emptyset$, because for $Z' \subseteq \overline{X}$ we have $\kappa \downarrow(Z') \geq \lambda(Y \cap Z') \geq |Y \cap Z'|$ by (6.B). By Lemma 6.1, we have $\kappa \downarrow(X \cap Z) < \kappa \downarrow(X)$. By symmetry, we also have $X \setminus Z = X \cap \overline{Z} \neq \emptyset$ and $\kappa \downarrow(X \setminus Z) < \kappa \downarrow(X)$. \square

6.2 A Canonical Family of Partitions

While so far, all our constructions work for general connectivity functions, in this section we need to restrict our attention to the cut rank function of a graph. In addition to Assumptions 5.1, which we still maintain, we make the following assumption.

Assumption 6.4. There is a graph G such that $\kappa = \rho_G$.

As in the previous subsection, let $X \subseteq A \downarrow$ such that $k_1 := \kappa \downarrow(X) \geq (3k+2)k$. Then

$$\text{rk}(M_{X \uparrow, \overline{X} \uparrow}) = k_1.$$

For every $W \subseteq A \downarrow = V(G) \downarrow$, we let $G[W \uparrow]$ be the induced subgraph of G with vertex set $W \uparrow$, and we let $\kappa \downarrow_W := \rho_{G[W \uparrow]} \downarrow$, where of course we contract only those C_i that are contained in $W \uparrow$. Observe that for every $Z \subseteq W$ we have

$$\kappa \downarrow_W(Z) = \text{rk}(M_{Z \uparrow, (W \setminus Z) \uparrow}).$$

By \overline{X}^ℓ we denote the set of all ℓ -tuples of elements of \overline{X} with mutually distinct entries. For every $\ell \geq 1$, we shall define an equivalence relation \equiv_X^ℓ on \overline{X}^ℓ with index (that is, number of equivalence classes) bounded in terms of k_1 and ℓ such that the following holds.

Lemma 6.5. *Let $\mathbf{w} = (w_1, \dots, w_\ell)$, $\mathbf{w}' = (w'_1, \dots, w'_\ell) \in \overline{X}^\ell$ such that $\mathbf{w} \equiv_X^\ell \mathbf{w}'$, and let $W := \{w_1, \dots, w_\ell\}$, $W' := \{w'_1, \dots, w'_\ell\}$.*

Let $Z \subseteq X \cup W$ and $Z' := (X \cap Z) \cup \{w'_i \mid i \in [\ell] \text{ such that } w_i \in Z\}$. Then

$$\kappa \downarrow_{X \cup W}(Z) = \kappa \downarrow_{X \cup W'}(Z').$$

Let us first consider the special case that no sets are contracted, that is, $m = -1$ (this is not a case that we actually need to consider, but it is helpful to explain the ideas). Then $A \downarrow = A = V(G)$ and $\kappa \downarrow = \kappa = \rho_G$. For $\mathbf{w} = (w_1, \dots, w_\ell)$, $\mathbf{w}' = (w'_1, \dots, w'_\ell) \in \overline{X}^\ell$, we let $\mathbf{w} \equiv_X^\ell \mathbf{w}'$ if for all $i \in [\ell]$ the columns of the matrix $M_{X, \overline{X}}$ indexed by w_i and w'_i are equal and for all $i, j \in [\ell]$ we have $w_i w_j \in E(G) \iff w'_i w'_j \in E(G)$. We can rephrase these two conditions as follows:

- (i) For all $i \in [\ell]$ the matrices $M_{X, \{w_i\}}$ and $M_{X, \{w'_i\}}$ are equal.
- (ii) For all $i \in [\ell]$ the matrices $M_{\{w_1, \dots, w_{i-1}, w_{i+1}, \dots, w_\ell\}, \{w_i\}}$ and $M_{\{w'_1, \dots, w'_{i-1}, w'_{i+1}, \dots, w'_\ell\}, \{w'_i\}}$ are equal.

Let $W := \{w_1, \dots, w_\ell\}$ and $W' := \{w'_1, \dots, w'_\ell\}$. For all $Z \subseteq X \cup W$ and $Z' := (X \cap Z) \cup \{w'_i \mid w_i \in Z\} \subseteq X \cup W'$, condition (i) implies that

$$M_{X \cap Z, W \setminus Z} = M_{X \cap Z', W' \setminus Z'}, \quad (6.E)$$

$$M_{X \setminus Z, W \cap Z} = M_{X \setminus Z', W' \cap Z'}. \quad (6.F)$$

Condition (ii) implies

$$M_{W \cap Z, W \setminus Z} = M_{W' \cap Z', W' \setminus Z'}. \quad (6.G)$$

Lemma 6.5 (in the special case) follows easily, because

$$\kappa \downarrow_{X \cup W}(Z) = \rho_{G[X \cup W]}(Z) = \text{rk}(M_{Z, (X \cup W) \setminus Z}) = \text{rk} \left(\begin{pmatrix} M_{X \cap Z, X \setminus Z} & M_{X \cap Z, W \setminus Z} \\ M_{W \cap Z, X \setminus Z} & M_{W \cap Z, W \setminus Z} \end{pmatrix} \right) \quad (6.H)$$

and

$$\kappa \downarrow_{X \cup W'}(Z') = \text{rk} \left(\begin{pmatrix} M_{X \cap Z', X \setminus Z'} & M_{X \cap Z', W' \setminus Z'} \\ M_{W' \cap Z', X \setminus Z'} & M_{W' \cap Z', W' \setminus Z'} \end{pmatrix} \right). \quad (6.I)$$

Since $X \cap Z' = X \cap Z$ and $X \setminus Z' = X \setminus Z$, equations (6.E), (6.F), and (6.G) imply that the matrices in the rightmost terms in (6.H) and (6.I) are equal.

Let us now turn to the general case. The situation is more difficult here because the sets C_i and hence the matrices involved in our argument above, in particular the matrices $M_{Z\uparrow, W\setminus Z\uparrow}$ may have unbounded size (in terms of k and ℓ). The crucial observation is that we can bound the size of the C_i in terms of $k \leq k_1$, exploiting the fact that $\rho_G(C_i) = \kappa\downarrow(\{c_i\}) < k$. To simplify the notation, we assume that $B = \{c_{m+1}, \dots, c_n\}$ for some $n \geq m$, and for $i = m+1, \dots, n$, we let $C_i := \{c_i\}$, so that actually $A = \bigcup_{i=0}^n C_i$ and $A\downarrow = \{c_0, \dots, c_n\}$.

Suppose that for some $i \in [m]$ there are distinct vertices $v, v' \in C_i$ such that for all $w \in V(G) \setminus C_i$ we have $vw \in E(G) \iff v'w \in E(G)$. Let $G' := G \setminus \{v'\}$ and $C'_i := C_i \setminus \{v'\}$ and $C'_j := C_j$ for $j \neq i$. Then contracting C'_1, \dots, C'_m in G' has the same effect as contracting C_1, \dots, C_m in G , that is,

$$\begin{aligned} A\downarrow &= V(G)\downarrow_{C_1, \dots, C_m} = V(G')\downarrow_{C'_1, \dots, C'_m}, \\ \kappa\downarrow &= \rho_G\downarrow_{C_1, \dots, C_m} = \rho_{G'}\downarrow_{C'_1, \dots, C'_m}. \end{aligned}$$

By repeating this construction we arrive at an induced subgraph $G'' \subseteq G$ and a partition $C''_1, \dots, C''_n \subseteq V(G'')$, where $C''_i \subseteq C_i$ for $0 \leq i \leq m$ and $C''_i = C_i = \{c_i\}$ for $m+1 \leq i \leq n$, such that

$$\begin{aligned} A\downarrow &= V(G'')\downarrow_{C''_1, \dots, C''_n}, \\ \kappa\downarrow &= \rho_{G''}\downarrow_{C''_1, \dots, C''_n}, \end{aligned}$$

and for all $i \in [n]$ and distinct $v, v' \in C''_i$ there is a $w \in V(G'') \setminus C''_i$ such that $vw \in E(G) \not\equiv v'w \in E(G)$. Observe that the construction of G'' and C''_1, \dots, C''_n from G and C_1, \dots, C_m is canonical and can be carried out in polynomial time. To simplify the notation, in the following we assume that $G'' = G$ and $C''_i = C_i$ for all $i \in [n]$.

Let $i \in [n]$. As for all distinct $v, v' \in C_i$ there is a $w \in \overline{C}_i$ such that $vw \in E(G) \not\equiv v'w \in E(G)$, the rows of the matrix M_{C_i, \overline{C}_i} are mutually distinct. As $k > \kappa\downarrow(\{c_i\}) = \rho_G(C_i) = \text{rk}(M_{C_i, \overline{C}_i})$, the matrix M_{C_i, \overline{C}_i} , being a matrix over \mathbb{F}_2 , has at most 2^{k-1} distinct rows. This implies that

$$|C_i| \leq 2^{k-1}.$$

Now we are ready to define the equivalence relation \equiv_X^ℓ . For this we let $\mathbf{w} = (w_1, \dots, w_\ell)$, $\mathbf{w}' = (w'_1, \dots, w'_\ell)' \in \overline{X}^\ell$. To simplify the notation, for every $i \in [\ell]$ we let $w_i\uparrow := \{w_i\}\uparrow$, that is, if $w_i = c_j$ then $w_i\uparrow = C_j$. Similarly, we let $w'_i\uparrow := \{w'_i\}\uparrow$. We let $\mathbf{w} \equiv_X^\ell \mathbf{w}'$ if for every $i \in [\ell]$ there are linear orders \leq_i of $w_i\uparrow$ and \leq'_i of $w'_i\uparrow$ such that the following two conditions are satisfied.

- (i) For all $i \in [\ell]$ the matrices $M_{X, w_i\uparrow}$ and $M_{X, w'_i\uparrow}$ are equal if the columns of the matrices are ordered according to the linear orders \leq_i and \leq'_i , respectively.
- (ii) For all $i \in [\ell]$ the matrices $M_{w_1\uparrow \cup \dots \cup w_{i-1}\uparrow \cup w_{i+1}\uparrow \cup \dots \cup w_\ell\uparrow, w_i\uparrow}$ and $M_{w'_1\uparrow \cup \dots \cup w'_{i-1}\uparrow \cup w'_{i+1}\uparrow \cup \dots \cup w'_\ell\uparrow, w'_i\uparrow}$ are equal if
 - (ii-a) the rows of the matrices are ordered lexicographically according to the natural order on the indices j of the w_j, w'_j and, within the sets $w_j\uparrow, w'_j\uparrow$, according to the linear orders \leq_j, \leq'_j , respectively;
 - (ii-b) columns of the matrices are ordered according to the linear orders \leq_i, \leq'_i , respectively.

Proof of Lemma 6.5. We have

$$\begin{aligned} \kappa\downarrow_{X \cup W}(Z) &= \rho_{G[(X \cup W)\uparrow]}(Z\uparrow) = \text{rk}(M_{Z\uparrow, ((X \cup W)\setminus Z)\uparrow}) \\ &= \text{rk} \left(\begin{pmatrix} M_{(X \cap Z)\uparrow, (X \setminus Z)\uparrow} & M_{(X \cap Z)\uparrow, (W \setminus Z)\uparrow} \\ M_{(W \cap Z)\uparrow, (X \setminus Z)\uparrow} & M_{(W \cap Z)\uparrow, (W \setminus Z)\uparrow} \end{pmatrix} \right). \end{aligned} \quad (6.J)$$

Similarly, for $Z' := (X \cap Z) \cup \{w'_i \mid w_i \in Z\}$,

$$\kappa \downarrow_{X \cup W}(Z) = \text{rk} \left(\begin{pmatrix} M_{(X \cap Z) \uparrow, (X \setminus Z) \uparrow} & M_{(X \cap Z) \uparrow, (W' \setminus Z') \uparrow} \\ M_{(W' \cap Z') \uparrow, (X \setminus Z) \uparrow} & M_{(W' \cap Z') \uparrow, (W' \setminus Z') \uparrow} \end{pmatrix} \right), \quad (6.K)$$

where we use the fact that $X \cap Z = X \cap Z'$ and $X \setminus Z = X \setminus Z'$. We may assume that in all these matrices the rows and columns indexed by entries of W, W' are ordered lexicographically according to the indices of the i and the orders \leq_i, \leq'_i as in (ii-a) above.

Observe that (i) implies

$$\begin{aligned} M_{(X \cap Z) \uparrow, (W \setminus Z) \uparrow} &= M_{(X \cap Z) \uparrow, (W' \setminus Z') \uparrow}, \\ M_{(W \cap Z) \uparrow, (X \setminus Z) \uparrow} &= M_{(W' \cap Z') \uparrow, (X \setminus Z) \uparrow}. \end{aligned}$$

Furthermore, (ii) implies that

$$M_{(W \cap Z) \uparrow, (W \setminus Z) \uparrow} = M_{(W' \cap Z') \uparrow, (W' \setminus Z') \uparrow}.$$

Thus the matrices in the rightmost terms in (6.J) and (6.K) are equal. \square

The following lemma collects further useful properties of the equivalence relation \equiv_X^ℓ .

Lemma 6.6. *Let $\ell \geq 1$*

- (1) *Given X , the equivalence relation \equiv_X^ℓ is canonical.*
- (2) *There is a $e_1 = e_1(k_1, \ell)$ (independent of $\kappa \downarrow$) such that the index of \equiv_X^ℓ is at most $e_1(k_1, \ell)$.*
- (3) *Given the graph G , the sets C_0, \dots, C_m , and the set X , the equivalence relation \equiv_X^ℓ can be computed in polynomial time (for fixed k_1 and ℓ).*

Proof. (1) and (3) are obvious from the construction.

(2) follows easily from the following two observations.

- For every $\mathbf{w} = (w_1, \dots, w_\ell) \in \overline{X}^\ell$ the set $\bigcup_{i=1}^\ell w_i \uparrow$ has at most $\ell \cdot 2^{k-1} \leq \ell \cdot 2^{k_1-1}$ elements.
- The matrix $M_{X \uparrow, \overline{X} \uparrow}$ has rank k_1 and thus at most 2^{k_1} different columns. \square

We are now ready to prove the main result of this section. Let

$$e_2(k_1) := \max\{e_1(k_1, \ell) \mid 1 \leq \ell \leq 2^{k_1}\}, \quad (6.L)$$

where e_1 is chosen according to Lemma 6.6.

Lemma 6.7. *Let $k_1 \geq (3k+2)k$ and $e_2 = e_2(k_1)$. Then for every $X \subseteq A \downarrow$ of order $\kappa \downarrow(X) = k_1$ there is a canonical family of $e \leq e_2$ partitions $X_1^{(i)}, X_2^{(i)}$ of X , for $1 \leq i \leq e$, such that $\kappa \downarrow(X_j^{(i)}) < k$ for $1 \leq i \leq e, j = 1, 2$.*

Furthermore, given X and oracle access to $\kappa \downarrow$, the family of partitions can be computed in polynomial time (for fixed k, k_1).

Proof. Let ℓ be the number of distinct columns of $M_{X \uparrow, \overline{X} \uparrow}$. Observe that $k \leq \ell \leq 2^{k_1}$. We call a set $W \subseteq \overline{X}$ *complete* if all columns of the matrix $M_{X \uparrow, \overline{X} \uparrow}$ already appear in the matrix $M_{X \uparrow, W \uparrow}$. A tuple $\mathbf{w} \in \overline{X}^\ell$ is *complete* if the set of its entries is complete. Observe that if \mathbf{w} is complete and $\mathbf{w}' \equiv_X^\ell \mathbf{w}$, then \mathbf{w}' is complete as well.

Let $\mathbf{w}^{(1)}, \dots, \mathbf{w}^{(e)}$ be a system of representatives of the \equiv_X^ℓ -equivalence classes consisting of complete tuples. Note that $e \leq e_2(k_1)$. For $1 \leq i \leq e$, let $W^{(i)}$ be the set of entries of $\mathbf{w}^{(i)}$ and $A \downarrow^{(i)} := X \cup W^{(i)}$ and $\kappa \downarrow^{(i)} := \kappa \downarrow_{A \downarrow^{(i)}}$. By Lemma 6.5, up to renaming of the elements,

the connectivity function $\kappa\downarrow^{(i)}$ only depends on the equivalence class of $\mathbf{w}^{(i)}$ and not on the choice of the specific tuple. Thus, up to renaming, the family of connectivity function $\kappa\downarrow^{(i)}$ is canonical.

Claim 1. Let $i \in [e]$. For each $Z \subseteq X$, we have $\kappa\downarrow^{(i)}(Z) = \kappa\downarrow(Z)$.

Proof. This follows from the completeness of $W^{(i)}$. \square

In particular, we have $\kappa\downarrow^{(i)}(X) = \kappa\downarrow(X) = k_1$. Now we apply the construction of Section 6.1 to $\kappa\downarrow^{(i)}$. We define $\lambda^{(i)} : 2^{W^{(i)}} \rightarrow \mathbb{N}$ by $\lambda^{(i)}(Y) := \kappa\downarrow_{\min}^{(i)}(Y, X)$ and let $\mathcal{M}(\lambda^{(i)})$ be the matroid induced by $\lambda^{(i)}$. Note that the order of the entries of the tuple $\mathbf{w}^{(i)}$ gives us a linear order $\leq^{(i)}$ on $W^{(i)}$. We let $Y^{(i)}$ be the lexicographically first subset of $W^{(i)} \setminus \{c_0\}$ of size $3k + 1$ that is independent in $\mathcal{M}(\lambda^{(i)})$. We let $Z_0^{(i)}$ be the lexicographically first subset of $Y^{(i)}$ such that for the leftmost minimum $(Z_0^{(i)}, Y^{(i)} \setminus Z_0^{(i)})$ -separation $Z^{(i)}$ we have

$$\kappa\downarrow(Z^{(i)}) < \min\{|Y^{(i)} \cap Z^{(i)}|, |Y^{(i)} \setminus Z^{(i)}|\}, \quad (6.M)$$

and we let $X_1^{(i)} := X \cap Z^{(i)}$ and $X_2^{(i)} := X \setminus Z^{(i)}$. By Lemma 6.3 we have $\kappa\downarrow^{(i)}(X_j^{(i)}) < \kappa\downarrow^{(i)}(X)$, and by Claim 1 this implies $\kappa\downarrow(X_j) < \kappa\downarrow(X)$.

Clearly, the construction is canonical and can be carried out in polynomial time. \square

Corollary 6.8. *For every k_1 there is a $c_1 = c_1(k, k_1)$ such that for every $X \subseteq A\downarrow$ of order $\kappa\downarrow(X) = k_1$ there is a canonical partial treelike decomposition (T_X, γ_X) with the following properties.*

- (i) T_X is a directed tree.
- (ii) $\gamma_X(r) = X$ for the root r of T_X .
- (iii) $\bigcup_{t \in L(T_X)} \gamma_X(t) = X$ (but the sets $\gamma_X(t)$ for the leaves $t \in L(T_X)$ are not necessarily disjoint).
- (iv) $\kappa\downarrow(\gamma_X(t)) \leq k_1$ for all $t \in V(T_X)$.
- (v) $\kappa\downarrow(\gamma_X(t)) < (3k + 2)k$ for all leaves $t \in L(T_X)$.
- (vi) T_X has at most $c_1(k, k_1)$ nodes.

Furthermore, given X and oracle access to $\kappa\downarrow$, the decomposition (T_X, γ_X) can be computed in polynomial time (for fixed k, k_1).

Proof. We define c_1 inductively by letting $c_1(k, k_1) = 1$ for all $k_1 < (3k + 2)k$ and

$$c_1(k, k_1) = 4e_2(k_1) \cdot c_1(k, k_1 - 1)$$

for every $k_1 \geq (3k + 2)k$.

We construct the decomposition T_X recursively as follows: we start with a root r and let $\gamma_X(r) := X$. If $\kappa\downarrow(X) < (3k + 2)k$, this is the whole decomposition. So suppose that $\kappa\downarrow(X) \geq (3k + 2)k$. Then we choose a canonical family of $e \leq e_1(k_1)$ partitions $X_1^{(i)}, X_2^{(i)}$ of X according to Lemma 6.7. For every $i \in [e]$ and $j = 1, 2$, let $(T_j^{(i)}, \gamma_j^{(i)})$ be the recursively constructed decomposition for $X_j^{(i)}$ (note that $\kappa\downarrow(X_j^{(i)}) < \kappa\downarrow(X) = k_1$).

To construct (T_X, γ_X) attach children $t^{(1)}, \dots, t^{(e)}$ to r and let $\gamma_X(t^{(i)}) := X$. For every i , we attach the trees $T_1^{(i)}, T_2^{(i)}$ to $t^{(i)}$ such that the root of $T_j^{(i)}$ becomes a child of $t^{(i)}$ in T_X . For all $t \in V(T_j^{(i)})$, we let $\gamma_X(t) := \gamma_j^{(i)}(t)$.

It is easy to see that this construction has the desired properties. \square

7 Constructing Canonical Treelike Decompositions

For the following lemma, we make Assumptions 5.1, 5.2, and 6.4.

Lemma 7.1. *There are $a_2 = a_2(k)$ and $g_1 = g_1(k)$ such that there is a treelike decomposition (T, γ) of $\kappa\downarrow$ with the following properties.*

- (i) T is a directed tree.
- (ii) T has at most n^{g_1} nodes, where $n := |A| = |V(G)|$.
- (iii) (T, γ) has width at most a_2 .
- (iv) $\beta(t) = \{c_0\}$ for the root r of T .
- (v) $|\gamma(t)| = 1$ for all leaves $t \in L(T)$.

Furthermore, given κ and C_0, \dots, C_m the construction of (T, γ) is canonical and can be carried out by a polynomial time algorithm (for fixed k).

Proof. We build the tree T recursively in a top-down fashion. We start the construction of T with the creation of a root r . We let $\gamma(r) := A\downarrow$. For every triple cover Q of the tangle \mathcal{T}_0 of size $|Q| \leq \theta(3k - 2)$ we create two fresh nodes s_Q, t_Q and edges (r, s_Q) and (s_Q, t_Q) . We let

$$\gamma(s_Q) := A\downarrow \setminus \{c_0\}, \quad (7.A)$$

$$\gamma(t_Q) := A\downarrow \setminus (Q^\vee \cup \{c_0\}), \quad (7.B)$$

where

$$Q^\vee := (B \cap Q) \cup \{c_i \mid 0 \leq i \leq m, C_i \cap Q \neq \emptyset\}.$$

(recall Assumptions 5.3). We call the nodes s_Q *cover nodes*.

Now suppose that t is some leaf of the current tree with $\gamma(t) = X$ for some X of size at least 2. Then some ancestor is the child t_Q of a cover node s_Q . We let Q be the triple cover associated with this node and Q^\vee its contraction. Thus now we may also make Assumption 5.3 and apply the results of Section 5. Note that $X \subseteq \gamma(t_Q) = A\downarrow \setminus (Q^\vee \cup \{c_0\})$.

If $\kappa\downarrow(X) < (3k + 2)k$, we apply Lemma 5.4 and obtain a partition X_1, \dots, X_p of X . We attach fresh nodes t_1, \dots, t_p as children to X and let $\gamma(t_i) := X_i$. In this case, we call t a *small node*. Depending on whether we apply Lemma 5.4(i) or (ii), we call t a small node of *type (i)* or *(ii)*.

If $\kappa\downarrow(X) \geq (3k + 2)k$, we apply Corollary 6.8 and obtain a partial treelike decomposition (T_X, γ_X) . We assume T_X to be disjoint from the tree we have constructed so far. We extend our tree by adding the tree T_X and identifying its root with t . For all $u \in V(T_X)$, we let $\gamma(u) := \gamma_X(u)$. In this case, we call t a *big node*.

The construction stops once $|\gamma(t)| = 1$ for all leaves. This guarantees property (v).

Property (i) is immediate from the construction. Property (iv) follows from (7.A), because the children of the root r are the cover nodes s_Q . The “bounded width” property (iii) follows from Lemma 5.4 and Corollary 6.7. This is clear for nodes whose degree is bounded in terms of k . In particular, this is the case for all small nodes of type (i) and all big nodes and all nodes that are introduced as part of a tree T_X at a big node t with $\gamma(t) = X$, because by Corollary 6.7(v) the tree T_X only has a bounded number of nodes. For small nodes of type (ii), it follows from Lemma 6.7(ii) that the width is bounded.

It remains to prove (ii), that is, bound the size of the tree. As there are only polynomially many (at most $n^{\theta(3k-2)+1}$) triple covers Q of \mathcal{T}_0 of size $|Q| \leq \theta(3k - 2)$, it suffices to prove a polynomial bound on the size of the trees $T^{(Q)}$ rooted in the nodes t_Q , the grandchildren of the root. So let us fix Q and only consider the tree $T' := T^{(Q)}$ from now on.

Next, we discard all internal nodes of the trees T_X associated with big nodes and draw direct edges from the root of each T_X to its leaves; this reduces the size of the tree only by a constant factor. Let T'' be the resulting tree. For each node $t \in V(T'')$ we define its *weight* to be

$$w(t) := |\gamma(t)|$$

Let

$$a_3 := \max\{a_1(k, (3k+2)k-1), 2(3k+2)k-2, (\theta(3k-2)+1)k\}.$$

Claim 1. For all nodes $t \in V(T'')$ we have $\kappa\downarrow(\gamma(t)) \leq a_3$.

Proof. Let $t \in V(T'')$. If $t = t_Q$ is the root of T'' , then

$$|A\downarrow \setminus \gamma(t)| = |Q \cup \{c_0\}| \leq \theta(3k-2) + 1$$

and thus

$$\kappa\downarrow(\gamma(t)) = \kappa\downarrow(A\downarrow \setminus \gamma(t)) \leq (\theta(3k-2) + 1)k.$$

Suppose next that t is a child of a small node s . Then $\kappa\downarrow(\gamma(s)) \leq (3k+2)k-1$. By Lemma 5.4, if s is a small node of type (i) then $\kappa\downarrow(\gamma(t)) \leq a_1(k, (3k+1)k-1)$ and if s is a small node of type (ii) then $\kappa\downarrow(\gamma(t)) \leq 2(3k+1)k-2$. \square

Let

$$\begin{aligned} b &:= \max\{b_1(k, a_3), c_1(k, a_3)\}, \\ q &:= (1 - 1/f_1(k, a_3)). \end{aligned}$$

Then our tree T'' has the following properties.

- (A) All internal nodes have at least two children.
- (B) All children of a big node are small nodes or leaves.
- (C) A big node has at most b children.
- (D) A small node of type (i) has at most b children.
- (E) For all children u of a small node t of type (i) it holds that $w(u) \leq q \cdot w(t)$.
- (F) A small node of type (ii) has at most one child u , called its *heavy child*, such that $w(u) > q \cdot w(t)$. For all other children u' it holds that $w(u') \leq q \cdot w(t)$.
- (G) If a small node of type (ii) has a heavy child, then this child is a small node as well.
- (H) For a small node (of either type) with children u_1, \dots, u_m it holds that

$$w(t) = \sum_{i=1}^m w(u_i).$$

- (I) For all leaves t it holds that $w(t) = 1$.

We re-structure the tree again. We inductively define the *small subtree* of a small node t to consist of t and the small subtrees of all small children of t . We let T''' be the tree obtained from T'' by contracting all small subtrees to their roots. That is, we delete all nodes that are small children of their parents, and add edges from every small node that remains to all children of leaves of its small subtree. Observe that $|T'''| \geq |T''|/2$. Thus it suffices to bound the size of T''' .

The tree T''' has the following properties.

- (J) All internal nodes have at least two children.
- (K) All children of a big node are small nodes or leaves.
- (L) A big node has at most b children.
- (M) For all children u of a small node t it holds that $w(u) \leq q \cdot w(t)$.
- (N) For a small node with children u_1, \dots, u_m it holds that

$$w(t) = \sum_{i=1}^m w(u_i).$$

- (O) For all leaves t it holds that $w(t) = 1$.

To estimate the size of T''' , we use the following arithmetic fact (more or less, Jensen's inequality).

Claim 2. Let $m \in \mathbb{N}$ and $c, q, x, x_1, \dots, x_m \in \mathbb{R}$ such that $c \geq 1$, $0 \leq q \leq 1$, $x > 0$, $0 \leq x_1, \dots, x_m \leq qx$, and $\sum_{i=1}^m x_i = x$. Then

$$\sum_{i=1}^m x_i^c \leq (q^c + (1-q)^c)x^c.$$

Proof. Without loss of generality we may assume that $q \geq 1/2$, because otherwise the claim for q follows from the claim for $q' = (1-q)$.

For all $x > 0$, $1/2 \leq q \leq 1$, $m \geq 2$, we let $\mathbb{D}(x, q, m)$ be the set of all $(x_1, \dots, x_m) \in \mathbb{R}^m$ such that $\sum_{i=1}^m x_i = x$ and $0 \leq x_i \leq qx$ for all $i \in [m]$. We shall prove that for all $(x_1, \dots, x_m) \in \mathbb{D}(x, q, m)$ we have

$$\sum_{i=1}^m x_i^c \leq (q^c + (1-q)^c)x^c. \quad (7.C)$$

The proof is by induction on m . For the base step $m = 2$, we just observe that the function $y \mapsto y^c + (x-y)^c$ is convex on the interval $[x-qx, qx]$ and thus assumes its maximum at the boundary.

For the inductive step $m \rightarrow m+1$, where $m \geq 2$, let $x > 0$ and $x_1, \dots, x_{m+1} \in \mathbb{D}(x, q, m+1)$. Then $(x_1, \dots, x_m + x_{m+1}) \in \mathbb{D}(x, q, m)$, and by the inductive hypothesis we have

$$\sum_{i=1}^{m-1} x_i^c + (x_m + x_{m+1})^c \leq (q^c + (1-q)^c)x^c.$$

By applying the inductive hypothesis $m = 2$ with $q = 1$, $x = x_m + x_{m+1}$, we see that $x_m^c + x_{m+1}^c \leq (x_m + x_{m+1})^c$. Thus

$$\sum_{i=1}^{m+1} x_i^c \leq \sum_{i=1}^{m-1} x_i^c + (x_m + x_{m+1})^c \leq (q^c + (1-q)^c)x^c. \quad \lrcorner$$

We choose $c \geq 1$ such that

$$(q^c + (1-q)^c) \leq \frac{1}{2(b+1)}, \quad (7.D)$$

For every node $t \in V(T''')$, let $s(t)$ be the number of nodes in the subtree of T''' rooted in t .

Claim 3.

- (1) For all big nodes $t \in V(T''')$ we have $s(t) \leq 1 + bw(t)^c$.
- (2) For all small nodes and leaves $t \in V(T''')$ we have $s(t) \leq w(t)^c$.

Proof. The proof is by induction on T''' . For the leaves t we have $s(t) = w(t) = 1$ and thus the claim is trivial.

For the inductive step, let t be a node with children u_1, \dots, u_m . Assume first that t is a small node. Let $x := w(t)$ and $x_i := w(u_i)$ for all $i \in [m]$. Then $1 \leq x_i \leq qx$ by (M) and $\sum_{i=1}^m x_i = x$ by (N). By the induction hypothesis, we have $s(u_i) \leq 1 + bx_i^c \leq (b+1)x_i^c$. Hence by Claim 1,

$$s(t) = 1 + \sum_{i=1}^m s(u_i) \leq 1 + (b+1) \sum_{i=1}^m x_i^c \leq 1 + (b+1)(q^c + (1-q)^c)x^c \leq 1 + \frac{w(t)^c}{2},$$

where the last inequality holds by (7.D). As $w(t) \geq 2$ by (J) and (O), we have $1 \leq w(t)^c/2$ and thus $s(t) \leq w(t)^c$.

Assume next that t is a big node. By (K), all the u_i are small nodes or leaves, and thus by the induction hypothesis we have $s(u_i) \leq w(u_i)^c \leq w(t)^c$. By (L) we have $m \leq b$. Thus

$$s(t) = 1 + \sum_{i=1}^m s(u_i) \leq 1 + bw(t)^c. \quad \lrcorner$$

Claim 2 implies that $|T'''| \leq 1 + bn^c$, which gives the desired polynomial bound.

It is clear from the construction that the decomposition (T, γ) can be computed in polynomial time in its input and the number of nodes of the tree T . \square

Theorem 7.2 (Canonical Decomposition Theorem). *Let $k \in \mathbb{N}$. Then there is an $a = a(k) \in \mathbb{N}$ and a polynomial time algorithm that, given a graph G of rank width at most k , computes a canonical treelike decomposition of ρ_G of width at most a .*

Proof. Let G be a graph of rank width k . We let $A := V(G)$ and $\kappa := \rho_G$. Let \mathfrak{T}_{\max} be the set of all maximal κ -tangles. By the Duality Theorem 2.2, we have $\mathfrak{T}_{\max} = \mathfrak{T}_{\max}^{\leq k}$, that is, all tangles in \mathfrak{T}_{\max} have order at most k .

Let $\mathcal{T}_{\text{root}} \in \mathfrak{T}_{\max}$ be arbitrary. In a first step, we show how to construct a decomposition that is canonical given $\mathcal{T}_{\text{root}}$. In a second step, we will show how to combine the decompositions for all $\mathcal{T}_{\text{root}} \in \mathfrak{T}_{\max}$.

Let (T_1, γ_1, τ_1) be the directed tree decomposition for \mathfrak{T}_{\max} computed by the algorithm of Theorem 4.4. Let r_1 be the root of T_1 .

Let $t_1 \in V(T_1)$, and let u_1, \dots, u_m be the children of t_1 in T_1 (possibly $m = 0$). Let $C_0^{(t_1)} := \overline{\gamma_1(t_1)}$ and $C_i^{(t_1)} := \gamma_1(u_i)$. We define $c_0^{(t_1)}, \dots, c_m^{(t_1)}$, $B^{(t_1)}$, $A \downarrow^{(t_1)}$ and $\kappa \downarrow^{(t_1)}$ as in Assumptions 5.1(5)–(7). Let $\mathcal{T}_0^{(t_1)} := \tau_1^{-1}(t_1)$ be the tangle associated with t_1 and $k_0^{(t_1)} := \text{ord}(\mathcal{T}_0^{(t_1)})$.

Observe that Assumptions 5.1, 5.2, and 6.4 are satisfied. This is trivial for Assumptions 5.1(1), (2), (5)–(7), 5.2(3) and 6.4. We have $\kappa(C_i) < k$, because $\overline{C}_0 = \gamma_1(t_1) \in \mathcal{T}_0^{(t_1)}$ and $C_i = \gamma_1(u_i) \in \tau_1^{-1}(u_i)$ by (DTD.2), and the order of all these tangles is at most $k = \text{bw}(\kappa)$. This implies Assumption 5.1(3). Assumption 5.1(4) follows from (DTD.2). Assumptions 5.2(1) and (2) follow from (DTD.1) and (DTD.2). We let $(T_2^{(t_1)}, \gamma_2^{(t_1)})$ be the decomposition computed by the algorithm of Lemma 7.1. Let $r_2^{(t_1)}$ be the root of $T^{(2)}$. By Lemma 7.1(iv) we have

$$(A) \quad \beta_2^{(t_1)}(r_2^{(t_1)}) = \{c_0^{(t_1)}\}$$

and thus

$$c_0^{(t_1)} \notin \gamma_2^{(t_1)}(t_2)$$

for all $t_2 \in V(T_2^{(t_1)}) \setminus \{r_2^{(t_1)}\}$.

We may assume without loss of generality that the decomposition $(T_2^{(t_1)}, \gamma_2^{(t_1)})$ has the following two properties:

- (B) $\beta_2^{(t_1)}(t_2) = \emptyset$ for all $t_2 \in V(T_2^{(t_1)}) \setminus (L(T_2^{(t_1)}) \cup \{r_2^{(t_1)}\})$;
- (C) $|\gamma_2^{(t_1)}(t_2)| = 1$ for all $t_2 \in L(T_2^{(t_1)})$.

If $(T_2^{(t_1)}, \gamma_2^{(t_1)})$ does not have these properties, we normalise it (see Lemma 4.1), but leave the bag of the root untouched.

Now we join the decompositions (T_1, γ_1) and $(T_2^{(t_1)}, \gamma_2^{(t_1)})$ for $t_1 \in V(T_1)$ into a new decomposition (D, γ) . The directed graph D is defined as follows.

- The node set $V(D)$ consist of all pairs (t_1, t_2) , where $t_1 \in V(T_1)$ and $t_2 \in V(T_2^{(t_1)})$.
- For $(t_1, t_2), (u_1, u_2) \in V(D)$, there is an edge from (t_1, t_2) to (u_1, u_2) in $E(D)$ if
 - either $t_1 = u_1$ and $(t_2, u_2) \in E(T_2^{(t_1)})$
 - or $(t_1, u_1) \in E(T_1)$ and t_2 is a leaf of $T_2^{(t_1)}$ with $\gamma_2^{(t_1)}(t_2)\uparrow = \gamma_1(u_1)$ and $u_2 = r_2^{(u_1)}$ is the root of $T_2^{(u_1)}$.

Thus we take disjoint copies of all trees $T_2^{(t_1)}$ for $t_1 \in V(T_1)$ and connect them along the edges of T_1 , connecting leaves of a tree $T_2^{(t_1)}$ to matching roots of trees $T_2^{(u_1)}$ for children u_1 of t_1 in T_1 .

Observe that for every $t_1 \in V(T_1)$ and for every leaf $t_2 \in V(T_2^{(t_1)})$, either (t_1, t_2) is a leaf of D (that is, a node of has out-degree 0) or it has precisely one child (that is, out-neighbour), which is the root of $T_2^{(u_1)}$ for some child u_1 of t_1 in T_1 . This follows from (C) and the fact that (T_1, γ_1) is a tree decomposition, and not just a tree like decomposition, so t_1 cannot have two distinct children u_1, u'_1 with the same cone $\gamma_1(u_1) = \gamma_1(u'_1)$.

We define $\gamma : V(D) \rightarrow A$ by

$$\gamma(t_1, t_2) := \begin{cases} \gamma_1(t_1) & \text{if } t_2 = r_2^{(t_1)}, \\ \gamma_2^{(t_1)}(t_2)\uparrow & \text{otherwise.} \end{cases}$$

Claim 1. (D, γ) is a treelike decomposition of κ .

Proof. As T_1 and the $T_2^{(t_1)}$ are all directed trees, D is a directed acyclic graph (not necessarily a tree, though), because a cycle in D would project to a cycle in T_1 or some $T_2^{(t_1)}$. This proves (TL.1).

To verify (TL.2), let $((t_1, t_2), (u_1, u_2)) \in E(D)$. If $t_1 = u_1$, then $(t_2, u_2) \in E(T_2^{(t_1)})$ and thus $\gamma_2^{(t_1)}(u_2) \subseteq \gamma_2^{(t_1)}(t_2)$ by (TL.2) for $(T_2^{(t_1)}, \gamma_2^{(t_1)})$. If $t_2 \neq r_2^{(t_1)}$ then

$$\gamma(u_1, u_2) = \gamma_2^{(t_1)}(u_2)\uparrow \subseteq \gamma_2^{(t_1)}(t_2)\uparrow = \gamma(t_1, t_2).$$

If $t_2 = r_2^{(t_1)}$ then

$$\gamma(u_1, u_2) = \gamma_2^{(t_1)}(u_2)\uparrow \subseteq (A \downarrow^{(t_1)} \setminus \{c_0^{(t_1)}\})\uparrow = \gamma_1(t_1) = \gamma(t_1, t_2),$$

where the inclusion $\gamma_2^{(t_1)}(u_2)\uparrow \subseteq (A \downarrow^{(t_1)} \setminus \{c_0^{(t_1)}\})\uparrow$ follows from (A).

If $(t_1, u_1) \in E(T_1)$, then

$$\gamma(t_1, t_2) = \gamma_2^{(t_1)}(t_2)\uparrow = \gamma_1(u_1) = \gamma(u_1, u_2).$$

To verify (TL.3), let $(t_1, t_2) \in V(D)$ and $(u_1, u_2), (v_1, v_2) \in N_+^D(t_1, t_2)$; without loss of generality distinct. Then $t_1 = u_1 = v_1$, because otherwise (t_1, t_2) would have out-degree one. Then u_2, v_2 are children of t_2 in $T_2^{(t_1)}$, and it follows from (TL.3) for $(T_2^{(t_1)}, \gamma_2^{(t_1)})$ that either $\gamma_2^{(t_1)}(u_2) = \gamma_2^{(t_1)}(v_2)$, which implies $\gamma(u_1, u_2) = \gamma(v_1, v_2)$, or $\gamma_2^{(t_1)}(u_2) \cap \gamma_2^{(t_1)}(v_2) = \emptyset$, which implies $\gamma(u_1, u_2) \cap \gamma(v_1, v_2) = \emptyset$.

Finally, (TL.4) holds because

$$\gamma(r_1, r_2^{(r_1)}) = \gamma_1(r_1) = A$$

by (TL.4) for (T_1, γ_1) . ┘

Claim 2. The width of (D, γ) is at most the maximum width of the $(T_2^{(t_1)}, \gamma_2^{(t_1)})$ for $t_1 \in V(T_1)$.

Proof. Let $(t_1, t_2) \in V(D)$.

If t_2 is a leaf of $T_2^{(t_1)}$ such that $\gamma_2^{(t_1)}(t_2)\uparrow = \gamma_1(u_1)$ for some child u_1 of t_1 in T_1 , then $(u_1, r_2^{(u_1)})$ is a child of (t_1, t_2) in D , and we have $\gamma(u_1, r_2^{(u_1)}) = \gamma_1(u_1) = \gamma(t_1, t_2)$. This implies $\beta(t_1, t_2) = \emptyset$. The width of (D, γ) at t is

$$\max\{\kappa(\emptyset), \kappa(\gamma_1(u_1))\},$$

which is also the width of $(T_2^{(t_1)}, \gamma_2^{(t_1)})$ at t_2 .

Otherwise, the children of (t_1, t_2) in D are (t_1, u_2) for all children u_2 of t_2 in $T_2^{(t_1)}$. Assume first that t_2 is not the root of $T_2^{(t_1)}$. Then $\beta(t_1, t_2) = \beta_2^{(t_1)}(t_2) = \emptyset$ by (B). For every subset U of the children of t_2 in $T_2^{(t_1)}$ we have

$$\kappa\left(\bigcup_{u_2 \in U} \gamma(t_1, u_2)\right) = \kappa\left(\bigcup_{u_2 \in U} \gamma_2^{(t_1)}\uparrow\right) = \kappa\downarrow^{(t_1)}\left(\bigcup_{u_2 \in U} \gamma_2^{(t_1)}\right).$$

It follows that the width of (D, γ) at t is equal to the width of $(T_2^{(t_1)}, \gamma_2^{(t_1)})$ at t_2 .

If $t_2 = r_2^{(t_1)}$, we can argue similarly. ┘

This completes the first (and main) step of our proof; the construction of a bounded-width treelike decomposition that is canonical given the “root tangle” $\mathcal{T}_{\text{root}}$.

To obtain a fully canonical decomposition, we carry out the construction above for every $\mathcal{T}_{\text{root}} \in \mathfrak{T}_{\text{max}}$. Say, we obtain decompositions $(D_1, \gamma_1), \dots, (D_m, \gamma_m)$. As there is only a linear number of maximal tangles, we can compute all these decompositions in polynomial time. Our final decomposition is simply the disjoint union of these decompositions $(D_1, \gamma_1), \dots, (D_m, \gamma_m)$. □

8 Matrices of Bounded Partition Rank

In this section we consider symmetric matrices $P \in \{0, 1, ?\}^{V \times V}$ with entries 0, 1, ? and row and column indices from a set V . We usually denote the entries of such a matrix P by p_{vw} , for $v, w \in V$, and we denote the row $(p_{vw} \mid w \in V)$ with index v by \mathbf{p}_v . We need no special notation for the columns and just refer to them via their indices $w \in V$.

If the ?-entries of such a matrix P form a block diagonal matrix, we call P a ?-block matrix. That is, $P \in \{0, 1, ?\}^{V \times V}$ is a ?-block matrix if it is symmetric and there are mutually disjoint

subsets $I_1, \dots, I_m \subseteq V$ such that $p_{vw} = ?$ if and only if there is a $j \in [m]$ such that $v, w \in I_j$. We call the sets I_1, \dots, I_m the *?-indices* of P , and we say that row \mathbf{p}_v has *?-index* I_j if $v \in I_j$ (similarly for columns). For disjoint subsets $B, C \subseteq \{I_1, \dots, I_m\}$, we let $P_{B,C}$ be the submatrix of P obtained by deleting all rows corresponding to indices that are not in B and deleting all columns corresponding to indices that are not in C . Note that $P_{B,C}$ is a $\{0, 1\}$ -matrix. We denote by $P_{B, \overline{B}}$ the matrix $P_{B, \{I_1, \dots, I_m\} \setminus B}$.

We say that the matrix P has *partition rank at most k* if for each partition of the family of *?-indices* into two parts B and \overline{B} , the submatrix $P_{B, \overline{B}}$ has rank at most k over \mathbb{F}_2 .

We are interested in *?-block matrices* and their partition rank because we can use them to describe the width of treelike decompositions of cut-rank functions. Let (D, γ) be a normal treelike decomposition of the cut rank function ρ_G of a graph G , and let $t \in V(D)$ be a node with children u_1, \dots, u_ℓ such that the children have pairwise disjoint cones. We define an *associated ?-block matrix* $P \in \{0, 1, ?\}^{V(G) \times V(G)}$ with entries p_{vw} defined as follows:

- if there is an $i \in [\ell]$ such that $v, w \in \gamma(u_i)$ then $p_{vw} = ?$;
- if $v, w \in \overline{\gamma(t)}$ then $p_{vw} = ?$;
- otherwise, if $vw \in E(G)$ then $p_{vw} = 1$ and if $vw \notin E(G)$ then $p_{vw} = 0$.

Note that the *?-indices* of P are the sets $\gamma(u_1), \dots, \gamma(u_\ell), \overline{\gamma(t)}$.

Lemma 8.1. *Let (D, γ) be a normal treelike decomposition of the cut rank function ρ_G of a graph G . Let $t \in V(D)$ be a node whose children have mutually disjoint cones, and let P be the *?-block matrix* associated with t . Then the partition rank of P is equal to the width of (D, γ) at t .*

Proof. Let u_1, \dots, u_ℓ be the children of t . Then

$$\mathcal{I} := \{\gamma(u_1), \dots, \gamma(u_\ell), \overline{\gamma(t)}\}.$$

is the set of *?-indices* of P . The width of the decomposition at node t is

$$\text{wd}(D, \gamma, t) = \max_{U \subseteq \{u_1, \dots, u_\ell\}} \rho_G\left(\bigcup_{u \in U} \gamma(u)\right) = \max_{B \subseteq \{\gamma(u_1), \dots, \gamma(u_\ell)\}} \text{rk}(P_{B, \mathcal{I} \setminus B}) = \max_{B \subseteq \mathcal{I}} \text{rk}(P_{B, \mathcal{I} \setminus B}),$$

where the first equality holds because (D, γ) is normal and thus $\beta(t) = \emptyset$ and the last equality holds because $\text{rk}(P_{B, \mathcal{I} \setminus B}) = \text{rk}(P_{\mathcal{I} \setminus B, B})$ for all $B \subseteq \mathcal{I}$. \square

An *extension* of a $\{0, 1, ?\}$ -vector is a $\{0, 1\}$ -vector obtained by replacing each *?*-entry by a 0 or a 1. That is, $\mathbf{x} = (x_v \mid v \in V) \in \{0, 1\}^V$ is an extension of $\mathbf{p} = (p_v \mid v \in V) \in \{0, 1, ?\}^V$ if $p_v \in \{0, 1\}$ implies $x_v = p_v$, for all $v \in V$. We say that two $\{0, 1, ?\}$ -vectors are *compatible* if they have a common extension. An *isomorphism* from a matrix $P \in \{0, 1, ?\}^{V \times V}$ to a matrix $P' \in \{0, 1, ?\}^{V' \times V'}$ is a bijective mapping $\varphi : V \rightarrow V'$ such that $p_{vw} = p'_{\varphi(v)\varphi(w)}$ for all $v, w \in V$, where as usual we denote the entries of P by p_{vw} and the entries of P' by $p'_{v'w'}$.

Let $P \in \{0, 1, ?\}^{V \times V}$ be a *?-block matrix*. An *extension set* for P is a set of vectors $\text{Ext} \subseteq \{0, 1\}^V$ such that every row in \mathbf{p}_v of P has an extension in Ext . If Ext is an extension set for P , then for every $v \in V$ we denote the set of all extensions of \mathbf{p}_v in Ext by $\text{Ext}(v)$. We call a construction that assigns an extension set to every *?-block matrix* *canonical* if for every two isomorphic *?-block matrices* $P \in \{0, 1, ?\}^{V \times V}$ and $P' \in \{0, 1, ?\}^{V' \times V'}$ and every isomorphism ψ from P to P' the following two conditions are satisfied.

- There is a bijection χ from Ext to Ext' such that $\chi(\text{Ext}(v)) = \text{Ext}'(\psi(v))$ for all $v \in V$.
- For every $\mathbf{x} = (x_v \mid v \in V) \in \text{Ext}$ with $\chi(\mathbf{x}) =: \mathbf{x}' = (x'_{v'} \mid v' \in V') \in \text{Ext}'$ and every $v \in V$ we have $x_v = x'_{\psi(v)}$.

Theorem 8.2. *Let $k \in \mathbb{N}$. Then there is an $e = e(k) \in \mathbb{N}$ and a polynomial time algorithm that, given a $?$ -block matrix $P \in \{0, 1, ?\}^{V \times V}$ of partition rank at most k , computes a canonical extension set $\text{Ext} \subseteq \{0, 1\}^V$ for P of size $|\text{Ext}| \leq e$.*

Proof. Let $P \in \{0, 1, ?\}^{V \times V}$ be a $?$ -block matrix, and let $\mathcal{I} = \{I_1, \dots, I_m\}$ be the set of $?$ -indices of P . Without loss of generality we assume that P neither has repeated rows nor repeated columns. Note that this implies that each $?$ -index I_j has size at most $|I_j| \leq 2^k$. Also note that the assumption implies that a row cannot be compatible with another row of the same index.

Let h be the function defined recursively by $h(0) := 1$ and $h(i+1) := 2^k \cdot h(i) + 4$ for all $i \geq 0$.

Claim 1. Let R be a set of rows of P that have mutually distinct $?$ -indices and are mutually incompatible. Then $|R| \leq h(k)$.

Proof. We say that two rows $\mathbf{p}_v, \mathbf{p}_{v'}$ are *compatible in column w* if either $p_{vw} = p_{v'w}$ or $p_{vw} = ?$ or $p_{v'w} = ?$. For a set $W \subseteq V$, rows $\mathbf{p}_v, \mathbf{p}_{v'}$ are *compatible in W* if they are compatible in all columns $w \in W$, and for a set $\mathcal{J} \subseteq \mathcal{I}$, rows $\mathbf{p}_v, \mathbf{p}_{v'}$ are *compatible in \mathcal{J}* if they are compatible in $\bigcup_{j \in \mathcal{J}} I_j$. Rows $\mathbf{p}_v, \mathbf{p}_{v'}$ are *incompatible in w, W , or \mathcal{J}* if they are not compatible in w, W , \mathcal{J} , respectively.

For a set R of rows of P , by $\mathcal{I}(R)$ we denote the set of $?$ -indices of the rows in R .

By induction on $i \geq 0$, we shall prove the following.

- (A) If $|R| > h(i)$ and $\mathcal{J} \subseteq \mathcal{I}$ such that the rows in R are mutually incompatible in \mathcal{J} , then there are sets $B \subseteq \mathcal{I}(R), C \subseteq \mathcal{J}$ such that $B \cap C = \emptyset$ and $\text{rk}(P_{B,C}) > i$.

As the partition rank of P is at most k , assertion (A) implies that $|R| \leq h(k)$, that is, the assertion of the claim.

For the base step $i = 0$, we observe that if $|R| > 1$ then there are at least two rows $\mathbf{p}, \mathbf{p}' \in R$ that are incompatible in some $J \in \mathcal{J}$. If we let B be the set consisting of the $?$ -indices of \mathbf{p}, \mathbf{p}' and $C := \{J\}$, we obtain a submatrix $P_{B,C}$ of rank at least 1.

For the inductive step $i \rightarrow i+1$, we assume that $|R| > h(i+1)$. We first prove that there are four $?$ -indices $J, J' \in \mathcal{I}(R), K, K' \in \mathcal{J}$ such that $\{J, J'\} \cap \{K, K'\} = \emptyset$ and $P_{\{J, J'\}, \{K, K'\}}$ has rank at least 2. Since $|R| > h(i+1) \geq 4$, we find distinct rows $\mathbf{p}_1, \dots, \mathbf{p}_4 \in R$ that all have at least one 1-entry in a column $w \in K$ for some $K \in \mathcal{J}$. For $i = 1, \dots, 4$, let $J_i \in \mathcal{I}(R)$ be the $?$ -index of row \mathbf{p}_i .

Case 1: There are distinct $i, j \in [4]$ and a $?$ -index $K \in \mathcal{J}$ such that rows \mathbf{p}_i and \mathbf{p}_j both have a 1-entry in a column $w \in K$.

Let $K' \in \mathcal{J}$ such that \mathbf{p}_i and \mathbf{p}_j are incompatible in K' . Then the matrix $P_{\{J_i, J_j\}, \{K, K'\}}$ has rank at least 2.

Case 2: For all distinct $i, j \in [4]$ there is no $?$ -index $K \in \mathcal{J}$ such that rows \mathbf{p}_i and \mathbf{p}_j both have a 1-entry in a column $w \in K$.

Let $K_1 \in \mathcal{J}$ such that \mathbf{p}_1 has a 1-entry in K_1 . There is at most one $i \in \{2, 3, 4\}$ such that $K_1 = J_i$. Without loss of generality, $J_2, J_3 \neq K_1$. Then $\mathbf{p}_2, \mathbf{p}_3$ only have 0-entries in K_1 . Let $K_2 \in \mathcal{J}$ such that $\mathbf{p}_2, \mathbf{p}_3$ are incompatible in K_2 . Without loss of generality we may assume that \mathbf{p}_2 has a 1-entry in K_2 . If $K_2 \neq J_1$, then the matrix $P_{\{J_1, J_2\}, \{K_1, K_2\}}$ has rank at least 2.

Suppose that $K_2 = J_1$. Let $K_3 \in \mathcal{J}$ such that \mathbf{p}_3 has a 1-entry in K_3 . By the assumption of Case 2, we have $K_3 \neq K_2 = J_1$. Then the matrix $P_{\{J_1, J_3\}, \{K_1, K_3\}}$ has rank at least 2.

In the following, we fix $J, J' \in \mathcal{I}(R), K, K' \in \mathcal{J}$ such that $\{J, J'\} \cap \{K, K'\} = \emptyset$ and $P_{\{J, J'\}, \{K, K'\}}$ has rank at least 2.

At most four rows in R have their $?$ -index in $\{J, J', K, K'\}$, because the rows in R have mutually distinct $?$ -indices. We define an equivalence relation \sim on the remaining rows in R by letting $\mathbf{p} \sim \mathbf{p}'$ if and only if the rows \mathbf{p} and \mathbf{p}' coincide in all columns in $J \cup J' \cup K \cup K'$. As the rank of the matrix $P_{\mathcal{I} \setminus \{J, J', K, K'\}, \{J, J', K, K'\}}$ is at most k , there are at most 2^k equivalence classes. Hence there is an equivalence class R' of size at least

$$\frac{|R| - 4}{2^k} > \frac{h(i+1) - 4}{2^k} = h(i).$$

Note that $\{J, J', K, K'\} \cap \mathcal{I}(R') = \emptyset$ by the definition of R' . Furthermore, the rows in R' are mutually incompatible in $\mathcal{J}' := \mathcal{J} \setminus \{J, J', K, K'\}$, because they are mutually incompatible, but coincide in all columns in $J \cup J' \cup K \cup K'$. By the induction hypothesis, there are disjoint sets $B' \subseteq \mathcal{I}(R')$ and $C' \subseteq \mathcal{J}'$ such that $B' \cap C' = \emptyset$ and $\text{rk}(P_{B', C'}) > i$. Let $B := B' \cup \{J, J'\}$ and $C := C' \cup \{K, K'\}$. Then $B \cap C = \emptyset$. We claim that $\text{rk}(P_{B, C}) > i + 1$. To see this, let \mathbf{p}, \mathbf{p}' be two rows of the matrix $P_{B, C}$ with $?$ -indices in $J \cup J'$ such that the projections of \mathbf{p}, \mathbf{p}' to the columns in $K \cup K'$ are linearly independent. As all rows in R' coincide on all columns in $K \cup K'$, either \mathbf{p} or \mathbf{p}' is not a linear combination of rows of the matrix $P_{B', C}$. This shows that $\text{rk}(P_{B, C}) > \text{rk}(P_{B', C}) \geq \text{rk}(P_{B', C'}) \geq i + 1$. This completes the proof of (A) and hence of Claim 1. \lrcorner

We say a row \mathbf{p}_v of P is *lonely* if there are less than $g(k) := (k+1) \cdot 2^k$ rows that are compatible with \mathbf{p}_v .

Claim 2. The number of lonely rows of P is at most $2^k g(k) h(k)$.

Proof. Let R be the set of all lonely rows of P , and let $r := |R|$. We shall prove that $r \leq 2^k g(k) h(k)$.

As $|I_j| \leq 2^k$ for all j , there is a subset $R_0 \subseteq R$ of size

$$r_0 := |R_0| \geq \frac{r}{2^k}.$$

such that the rows $\mathbf{p} \in R_0$ have mutually distinct $?$ -indices. As the rows $\mathbf{p} \in R_0 \subseteq R$ are lonely, each such row is compatible with less than $g(k)$ rows. Thus there is a subset $R_1 \subseteq R_0$ of size

$$r_1 := |R_1| \geq \frac{r_0}{g(k)} \geq \frac{r}{2^k g(k)}$$

such that the rows in R_1 are mutually incompatible.

It follows from Claim 1 that $r_1 \leq h(k)$. Thus $r \leq 2^k g(k) h(k)$. \lrcorner

We call a $\{0, 1\}$ -vector a *supported extension* if it is an extension of at least $k+2$ rows of P .

Claim 3. The number of supported extensions is at most 2^k .

Proof. Let $\ell := 2^k + 1$. Suppose for contradiction that $\mathbf{x}_1, \dots, \mathbf{x}_\ell$ is a sequence of pairwise distinct supported extensions. Let X be the $[2^k + 1] \times V$ -matrix whose rows are the vectors \mathbf{x}_i . The rank of this matrix is at least $k+1$, because a matrix of rank at most k over \mathbb{F}_2 has at most 2^k distinct rows. Let $w_1, \dots, w_{k+1} \in V$ be column indices such that the columns of X indexed by the w_j are linearly independent, and let X' be submatrix of X obtained by deleting all columns except those with indices w_1, \dots, w_{k+1} . Then the rank of X' is $k+1$.

For all $j \in [k+1]$, let $K_j \in \mathcal{I}$ be the $?$ -index of column w_j of P . Let

$$C := \{K_1, \dots, K_{k+1}\}.$$

Let $i \in [\ell]$. The vector \mathbf{x}_i is an extension of distinct rows $\mathbf{p}_i^1, \dots, \mathbf{p}_i^{k+2}$ of P . These rows must have mutually distinct $?$ -indices, because distinct rows with the same $?$ -index cannot have a common extension. As $|C| = k+1$, there is a $j \in [k+2]$ such that the $?$ -index of \mathbf{p}_i^j is not in C .

We let $\mathbf{p}_i := \mathbf{p}_i^j$. Then \mathbf{p}_i and its extension \mathbf{x}_i coincide on all columns in C . Let J_i the $?$ -index of \mathbf{p}_i . Then $J_i \notin C$.

Now let $B := \{J_1, \dots, J_\ell\}$. Then $B \cap C = \emptyset$. Then matrix $P_{B,C}$ contains the rank- $(k+1)$ matrix X' as a submatrix, obtained by deleting all rows except the rows \mathbf{p}_i and all columns except those with indices w_1, \dots, w_{k+1} . This contradicts the partition rank of P being at most k . \perp

Let $\text{Ext} \subseteq \{0, 1\}^V$ be the union of the set of all vectors \mathbf{x} that are supported extensions and the set of all vectors \mathbf{y} that are extensions of lonely rows of P . Every row \mathbf{p} of P has at most 2^{2^k} extensions, because its $?$ -index has size at most 2^k . Thus, by Claims 2 and 3,

$$|\text{Ext}| \leq 2^k + 2^{2^k+k} \cdot g(k) \cdot h(k).$$

We claim that every row \mathbf{p} of P has an extension in Ext . This is obvious if \mathbf{p} is lonely. Suppose \mathbf{p} is not lonely. Let R be the set of all rows of P that are compatible with \mathbf{p} . Then $|R| \geq g(k) = (k+1) \cdot 2^k$. There is a subset $R' \subseteq R$ of size

$$|R'| = \frac{|R|}{2^k} \geq k+1$$

such that all rows in R' coincide on the columns in the $?$ -index I of \mathbf{p} . We extend \mathbf{p} to a vector $\mathbf{x} \in \{0, 1\}^V$ that coincides with the vectors in $|R'|$. Then \mathbf{x} is a common extension of \mathbf{p} and all vectors in R' . Thus \mathbf{x} is a supported extension and hence in Ext .

To finish the proof, it suffices now to note that the construction of Ext is canonical. \square

9 Computing the Automorphism Groups

We use various standard algorithms for permutation groups. Recall that a permutation group Γ that permutes elements in some set V can be succinctly represented by a generating set. For a set $\{g_1, \dots, g_t\}$ of permutations on V the group generated by the set is denoted by $\langle g_1, \dots, g_t \rangle$ or by $\langle \{g_1, \dots, g_t\} \rangle$. For sets V' and V , slightly abusing terminology, we call a set of bijections Λ from V' to V a (V, V') -coset, or just a coset if V and V' are clear from the context, if there is a bijection σ in Λ and a permutation group Γ on V' such that $\sigma\Gamma = \Lambda$.⁴ We also regard the empty set as a coset.

Our typical example of a (V, V') -coset is the set $\text{Iso}(G, G')$ of all isomorphism from a graph G with vertex set V to a graph G' with vertex set V' . Indeed, if $\text{Iso}(G, G')$ is nonempty then for every isomorphism φ from G to G' we have $\text{Iso}(G, G') = \varphi \text{Aut}(G')$, where $\text{Aut}(G')$ denote the automorphism group of G' .

It is important to note that if $\sigma\Gamma$ is a coset, then for all $\sigma' \in \sigma\Gamma$ it holds that

$$\sigma^{-1}\sigma' \in \Gamma \tag{9.A}$$

and

$$\sigma\Gamma = \sigma'\Gamma. \tag{9.B}$$

Indeed, writing $\sigma' = \sigma g'$ for some $g' \in \Gamma$ we have $\sigma^{-1}\sigma' = g' \in \Gamma$ and $\sigma g = \sigma'(g')^{-1}g \in \sigma'\Gamma$ and $\sigma'g = \sigma g'g \in \sigma\Gamma$ for all $g \in \Gamma$.

If we want to indicate that Λ is a subcoset of a coset Λ' (rather than just a subset), we write $\Lambda \leq \Lambda'$ (instead of $\Lambda \subseteq \Lambda'$). We will always assume that nonempty cosets are succinctly represented by one explicit bijection σ in Λ and a generating set for the permutation group Γ . There is a standard polynomial time algorithm that computes a succinct representation for a given coset. Permutations from the group Γ will be applied from left to right. That is,

⁴To relate this to the standard group theoretic notion of coset, note that if $V = V'$, then a (V, V') -coset is a left coset of a subgroup of the symmetric group on V .

for $g_1, g_2 \in \Gamma$, the element $g_1 g_2$ is the permutation that first applies g_1 and then applies g_2 . Note however that for maps φ we denote by $\varphi(x)$ the image of x under φ . It will be apparent from context which of the two notations is used. For more details on the algorithmic theory of permutation groups we refer to [33].

Let G and G' be graphs with vertex sets V, V' , respectively, and let (D, γ) and (D', γ') be directed decompositions of V and V' , respectively. An isomorphism φ from G to G' is said to *respect* (D, γ) and (D', γ') if there is an isomorphism $\hat{\varphi}$ from D to D' such that $\gamma'(\hat{\varphi}(t)) = \varphi(\gamma(t))$ for all $t \in V(D)$. We sometimes say that $\hat{\varphi}$ is an *isomorphism from (D, γ) to (D', γ') extending φ* . We denote the coset of all isomorphism from G to G' that respect (D, γ) and (D', γ') by $\text{Iso}(G_{D, \gamma}, G'_{D', \gamma'})$. Note that

$$\text{Iso}(G_{D, \gamma}, G'_{D', \gamma'}) \leq \text{Iso}(G, G').$$

Before we proceed to our isomorphism test for graphs of bounded rank width, we prove two simple lemmas about cosets. A *least upper bound* for two (V, V') -cosets Λ_1, Λ_2 (in the lattice of all (V, V') -cosets) is a (V, V') -coset Λ such that $\Lambda_1, \Lambda_2 \leq \Lambda$ and $\Lambda \leq \Lambda'$ for all cosets Λ' with $\Lambda_1, \Lambda_2 \leq \Lambda'$. Clearly, if a least upper bound exists then it is unique. The next lemma shows that indeed least upper bounds exist (this is easy to see) and can be computed in polynomial time.

Lemma 9.1. *There is a polynomial time algorithm that, given two (V, V') -cosets Λ_1, Λ_2 , computes the least upper bound of Λ_1 and Λ_2 .*

Of course we assume that both the input cosets Λ_1, Λ_2 and the output cosets are represented succinctly via generators.

Proof. Without loss of generality we assume that Λ_1 and Λ_2 are nonempty. Suppose that $\Lambda_i = \sigma_i \Gamma_i$, where $\sigma_i : V \rightarrow V'$ and $\Gamma_i \leq \text{Sym}(V')$ is represented by generators g_{i1}, \dots, g_{it_i} . We let $\sigma := \sigma_1$ and

$$\Gamma := \left\langle \{g_{1j} \mid j \in [t_1]\} \cup \{\sigma_1^{-1} \sigma_2 g_{2j} \mid j \in [t_2]\} \cup \{\sigma_1^{-1} \sigma_2 g_{2j} \sigma_2^{-1} \sigma_1 \mid j \in [t_2]\} \right\rangle \quad (9.C)$$

Let $\Lambda := \sigma \Gamma$. Then $\Lambda_1 \leq \Lambda$, because $\sigma_1 = \sigma$ and $\Gamma_1 = \langle \{g_{1j} \mid j \in [t_1]\} \rangle \leq \Gamma$. To see that $\Lambda_2 \leq \Lambda$, let $\sigma_2 g_2 \in \Lambda_2$, where $g_2 = g_{2j_1} g_{2j_2} \cdots g_{2j_\ell} \in \Gamma_2$. Then

$$\sigma_2 g_2 = \underbrace{\sigma_1^{-1} \sigma_2 g_{2j_1} \sigma_2^{-1} \sigma_1}_{=\sigma} (\underbrace{\sigma_1^{-1} \sigma_2 g_{2j_2} \sigma_2^{-1} \sigma_1 \cdots \sigma_1^{-1} \sigma_2 g_{2j_\ell} \sigma_2^{-1} \sigma_1}_{\in \Gamma} \sigma_1^{-1} \sigma_2 g_{2j_{\ell-1}}) \in \sigma \Gamma = \Lambda.$$

Let $\Lambda' = \sigma' \Gamma'$ be a coset with $\Lambda_1, \Lambda_2 \leq \Lambda'$. We need to prove that $\Lambda \leq \Lambda'$.

As $\sigma_i \in \Lambda_i \leq \Lambda'$, by (9.B) we have $\Lambda' = \sigma_i \Gamma'$, and thus by (9.A) we have $\sigma_i^{-1} \sigma_{3-i} \in \Gamma'$. We observe next that $\Gamma_i \leq \Gamma'$. Indeed, let $g \in \Gamma_i$. Then $\sigma_i g \in \Lambda_i \leq \Lambda = \sigma_i \Gamma'$ and thus $g = \sigma_i^{-1} \sigma_i g \in \Gamma'$.

Hence all generators of Γ are in Γ' . It follows that $\Gamma \leq \Gamma'$ and thus $\Lambda = \sigma_1 \Gamma \leq \sigma_1 \Gamma' = \Lambda'$.

Clearly, σ and the system of generators for Γ in (9.C) can be computed in polynomial time given the σ_i and the generators g_{ij} . \square

Lemma 9.2. *Let Λ be a (V, V') -coset. Let $W \subseteq V$ and $\varphi : W \rightarrow V'$ an injective mapping. Then the set*

$$\Lambda_{[\varphi]} := \{\psi \in \Lambda \mid \psi(w) = \varphi(w) \text{ for all } w \in W\}$$

is a subcoset of Λ . Furthermore, there is a polynomial time algorithm that, given Λ, W , an φ , computes $\Lambda_{[\varphi]}$.

Proof. Without loss of generality we assume that $\Lambda = \sigma\Gamma \neq \emptyset$. If there is a $\psi \in \Lambda$ such that $\psi(w) = \varphi(w)$ for all $w \in W$, then $\Lambda_{[\varphi]} = \psi\Gamma_{\varphi(W)}$, where $\Gamma_{\varphi(W)} \leq \Gamma$ is the pointwise stabiliser of $\varphi(W)$ in Γ . Otherwise, $\Lambda_{[\varphi]} = \emptyset$. Thus $\Lambda_{[\varphi]}$ is a coset.

For the algorithmic claim, it suffices to show that for given $w \in V$ and $w' \in V'$ it is possible to compute the subcoset $\Lambda_{w,w'}$ of those elements in $\Lambda = \sigma\Gamma$ that map w to w' . For this we find an element $g \in \Gamma$ that maps $\sigma(w)$ to v' using the standard orbit algorithm (see [33]). If no such element exists then $\Lambda_{w,w'} = \emptyset$. Otherwise $\Lambda_{w,w'} = \sigma g\Gamma_{w'}$ where $\Gamma_{w'}$ is the stabiliser of w' in Γ . The stabiliser of an element can also be computed in polynomial time (see [33]). \square

The main technical result of this section is the following theorem. Combined with the Canonical Decomposition Theorem (Theorem 7.2), it yields a polynomial time isomorphism test for graphs of bounded rank width (Theorem 9.5).

Theorem 9.3. *For every $k \in \mathbb{N}$ there is a polynomial time algorithm that, given graphs G, G' and normal treelike decompositions $(D, \gamma), (D', \gamma')$ of $\rho_G, \rho_{G'}$, respectively, of width at most k , computes a coset Λ such that*

$$\text{Iso}(G_{D,\gamma}, G'_{D',\gamma'}) \leq \Lambda \leq \text{Iso}(G, G').$$

Note that Λ may be empty if $\text{Iso}(G_{D,\gamma}, G'_{D',\gamma'})$ is empty, that is, if there is no isomorphism from G to G' that respects the treelike decompositions. It is worth noting that it is isomorphism-hard to compute $\text{Iso}(G_{D,\gamma}, G'_{D',\gamma'})$ even if G, G' are graphs with no edges and $(D, \gamma), (D', \gamma')$ treelike decompositions of width 0. Thus computing a coset that is sandwiched between $\text{Iso}(G_{D,\gamma}, G'_{D',\gamma'})$ and $\text{Iso}(G, G')$ is a crucial trick.

In the following, let us fix k and G, G' and $(D, \gamma), (D', \gamma')$. Let $V := V(G)$ and $V' := V(G')$.

Let $t \in V(D)$. We shall define a graph G_t that represents the induced subgraph $G[\gamma(t)]$ as well as an “abstraction” of the edges from $\gamma(t)$ to $\overline{\gamma(t)} = V \setminus \gamma(t)$. Let $W_t \subseteq \{0, 1\}^{\gamma(t)}$ be the set of rows that appear in the matrix $M_{\gamma(t), \gamma(t)}$. Since the width of $(D, \gamma(t))$ is at most k , the rank of the matrix $M_{\gamma(t), \gamma(t)}$ is at most k , and thus the set W_t has size at most 2^k . We may view the elements $\mathbf{w} = (w_v \mid v \in \gamma(t))$ in W_t as “types”, or equivalence classes of vertices $w \in \overline{\gamma(t)}$, where two vertices w, w' have the same type, or are equivalent, if they have the same adjacencies with the vertices in $\gamma(t)$. The entries of the vector \mathbf{w} are these adjacencies; $w_v = 1$ means that all vectors of this type are adjacent to v and $w_v = 0$ means that they are not adjacent.

Now we are ready to define the graph G_t . The vertex set is

$$V(G_t) := \gamma(t) \cup W_t,$$

and the edge set is

$$\begin{aligned} E(G_t) := & \{vv' \mid v, v' \in \gamma(t) \text{ such that } vv' \in E(G)\} \\ & \cup \{v\mathbf{w} \mid v \in \gamma(t) \text{ and } \mathbf{w} = (w_v \mid v' \in \gamma(t)) \in W_t \text{ such that } w_v = 1\}. \end{aligned}$$

Thus W_t is an independent set in G_t , and $G_t[\gamma(t)] = G[\gamma(t)]$. We colour the graph G_t so that the vertices in W_t are coloured red and all other vertices are coloured blue. We let D_t be the induced subgraph of D whose vertex set consist of all vertices that are reachable from t in D , and we let γ_t be the restriction of γ to $V(D_t)$. Then (D_t, γ_t) is a normal treelike decomposition of $V(G_t) \setminus W_t$. (We may also view it as a partial treelike decomposition of $V(G_t)$ or even $V(G)$; this does not matter, as we are not interested in the width of this decomposition.) Note that if r is the unique root of D , which exists by (NTL.4), then we have $G_r = G$ and $(D_r, \gamma_r) = (D, \gamma)$.

We define sets $W'_t \subseteq \{0, 1\}^{\gamma'(t)}$, graphs G'_t , and decompositions (D'_t, γ'_t) analogously for all nodes $t \in V(D')$.

Recall that our goal is to compute a coset Λ such that

$$\text{Iso}(G_{D,\gamma}, G'_{D',\gamma'}) \leq \Lambda \leq \text{Iso}(G, G')$$

We do this by a dynamic programming algorithm that processes the nodes of D in a bottom-up manner (starting from the leaves). The next lemma describes the inductive step.

Lemma 9.4. *There is a polynomial time algorithm that, given $G, G', (D, \gamma), (D', \gamma')$ as above and in addition*

- nodes $t \in V(D)$ and $t' \in V(D')$;
- for all $u \in N_+^D(t)$ and $u' \in N_+^{D'}(t')$ a coset $\Lambda(u, u')$ satisfying

$$\text{Iso}((G_u)_{D_u, \gamma_u}, (G'_{u'})_{D'_{u'}, \gamma'_{u'}}) \leq \Lambda(u, u') \leq \text{Iso}(G_u, G'_{u'}), \quad (9.D)$$

computes a coset Λ such that

$$\text{Iso}((G_t)_{D_t, \gamma_t}, (G'_{t'})_{D'_{t'}, \gamma'_{t'}}) \leq \Lambda \leq \text{Iso}(G_t, G'_{t'}).$$

Proof. Let $t \in V(D)$ and $t' \in V(D')$, and for all $u \in N_+^D(t)$ and $u' \in N_+^{D'}(t')$, let $\Lambda(u, u')$ be a coset satisfying (9.D).

We split the lengthy proof into several steps.

Step 1: The easy cases.

Suppose first that t is a leaf of D . As the decomposition (D, γ) is normal, by (NTL.2) we have $|\gamma(t)| = 1$. In this case, we can compute $\text{Iso}(G_t, G'_{t'})$ by brute force and let $\Lambda := \text{Iso}(G_t, G'_{t'})$.

In the following, we assume that t has at least one child. By (NTL.3) either all children of t have the same cone or the children have mutually disjoint cones.

Suppose that all children of t have the same cone. As $\beta(t) = \emptyset$, this implies that for all $u \in N_+^D(t)$ we have $\gamma(u) = \gamma(t)$. This implies $W_u = W_t$ and $G_u = G_t$. (But note that $(D_u, \gamma_u) \neq (D_t, \gamma_t)$!)

If there is a child u' of t' with $\gamma'(u') \neq \gamma'(t')$, then $\text{Iso}((G_t)_{D_t, \gamma_t}, (G'_{u'})_{D'_{u'}, \gamma'_{u'}}) = \emptyset$, and we can let $\Lambda := \emptyset$. So we assume that for all $u' \in N_+^{D'}(t')$ we have $\gamma'(u') = \gamma'(t')$. Again, this implies $W_{u'} = W_{t'}$ and $G'_{u'} = G'_{t'}$.

Hence for all $u \in N_+^D(t)$, $u' \in N_+^{D'}(t')$ we have $\text{Iso}(G_u, G'_{u'}) = \text{Iso}(G_t, G'_{t'})$ and thus

$$\text{Iso}((G_u)_{D_u, \gamma_u}, (G'_{u'})_{D'_{u'}, \gamma'_{u'}}) \leq \Lambda(u, u') \leq \text{Iso}(G_t, G'_{t'}) \quad (9.E)$$

We let Λ be the least upper bound of the cosets $\Lambda(u, u')$ for $u \in N_+^D(t)$, $u' \in N_+^{D'}(t')$, which we can compute in polynomial time by Lemma 9.1. Then by (9.E), for all $u \in N_+^D(t)$, $u' \in N_+^{D'}(t')$ we have

$$\text{Iso}((G_u)_{D_u, \gamma_u}, (G'_{u'})_{D'_{u'}, \gamma'_{u'}}) \leq \Lambda \leq \text{Iso}(G_t, G'_{t'})$$

It remains to prove that $\text{Iso}((G_t)_{D_t, \gamma_t}, (G'_{t'})_{D'_{t'}, \gamma'_{t'}}) \leq \Lambda$. Let $\psi \in \text{Iso}((G_t)_{D_t, \gamma_t}, (G'_{t'})_{D'_{t'}, \gamma'_{t'}})$, and let $\hat{\psi}$ be an isomorphism from (D_t, γ_t) to $(D'_{t'}, \gamma'_{t'})$ extending ψ . Let $u \in N_+^D(t)$ and $u' := \hat{\psi}(u)$. Then the restriction of $\hat{\psi}$ to D_u is an isomorphism from (D_u, γ_u) to $(D'_{u'}, \gamma'_{u'})$ that extends the isomorphism ψ from $G_u = G_t$ to $G'_{u'} = G'_{t'}$. Hence $\psi \in \text{Iso}((G_u)_{D_u, \gamma_u}, (G'_{u'})_{D'_{u'}, \gamma'_{u'}}) \leq \Lambda$.

Step 2: Fixing the outside.

In the following, we assume that the children of t have mutually disjoint cones. If the children of t' have the same cone, then $\text{Iso}((G_t)_{D_t, \gamma_t}, (G'_{t'})_{D'_{t'}, \gamma'_{t'}}) = \emptyset$, and we can let $\Lambda := \emptyset$. So we assume that the children of t' have mutually disjoint cones as well.

The colouring of the graphs G_t and $G'_{t'}$ forces each isomorphism $\psi \in \text{Iso}(G_t, G'_{t'})$ to map W_t bijectively to $W'_{t'}$. For each bijection $\varphi : W_t \rightarrow W'_{t'}$, we shall compute a coset Λ_φ such that

$$\text{Iso}((G_t)_{D_t, \gamma_t}, (G'_{t'})_{D'_{t'}, \gamma'_{t'}})_{[\varphi]} \leq \Lambda_\varphi \leq \text{Iso}(G_t, G'_{t'});$$

recall that $\text{Iso}((G_t)_{D_t, \gamma_t}, (G'_{t'})_{D'_{t'}, \gamma'_{t'}})_{[\varphi]}$ denotes the subcoset of $\text{Iso}((G_t)_{D_t, \gamma_t}, (G'_{t'})_{D'_{t'}, \gamma'_{t'}})$ consisting of all isomorphisms ψ whose restriction to W_t is φ . Then we let Λ be the least upper bound of the cosets Λ_φ (see Lemma 9.1). As $|W_t|, |W'_{t'}| \leq 2^k$, there are at most $2^k!$ bijections $\varphi : W_t \rightarrow W'_{t'}$, and if we can compute each Λ_φ in polynomial time, we can also compute Λ in polynomial time (for fixed k).

From now on, we fix a bijection $\varphi : W_t \rightarrow W'_{t'}$.

Let $u \in N_+^D(t)$ and $u' \in N_+^{D'}(t')$. Recall that $\gamma(u) \subseteq \gamma(t)$. The u -projection of a vector $\mathbf{w} = (w_v \mid v \in \gamma(t)) \in \{0, 1\}^{\gamma(t)}$ is the vector $\mathbf{w}|_u = (w_v \mid v \in \gamma(u)) \in \{0, 1\}^{\gamma(u)}$. We observe that for every vector $\mathbf{w} \in W_t$ we have $\mathbf{w}|_u \in W_u$. To see this, recall that \mathbf{w} appears as a row in the matrix $M_{\overline{\gamma(t)}, \gamma(t)}$. Let $w \in \overline{\gamma(t)}$ be the index of this row. Then $w \in \overline{\gamma(u)} \supseteq \overline{\gamma(t)}$, and hence there is also a row with index w in the matrix $M_{\overline{\gamma(u)}, \gamma(u)}$, and this row is precisely $\mathbf{w}|_u$. Similarly, for $u' \in N_+^{D'}(t')$, every $\mathbf{w}' \in W'_{t'}$ has a u' -projection $\mathbf{w}'|_{u'}$ in $W'_{u'}$.

Now let $\psi \in \Lambda(u, u')$. Then ψ induces a bijection from W_u to $W'_{u'}$. We say that ψ agrees with φ if for every $\mathbf{w} = (w_v \mid v \in \gamma(t)) \in W_t$ it holds that $\psi(\mathbf{w}|_u) = \varphi(\mathbf{w})|_{u'}$. Observe that with $\mathbf{w}' = (w'_{v'} \mid v' \in \gamma'(t')) := \varphi(\mathbf{w})$, this implies $w_v = w'_{\psi(v)}$ for all $v \in \gamma(u)$. Indeed, for $v \in \gamma(u)$ we have

$$w_v = 1 \iff v\mathbf{w}|_u \in E(G_u) \iff \psi(v)\mathbf{w}'|_{u'} \in E(G'_{u'}) \iff w'_{\psi(v)} = 1,$$

because ψ is an isomorphism from G_u to $G'_{u'}$. We denote the subcoset of all $\psi \in \Lambda(u, u')$ that agree with φ by $\Lambda(u, u')_{[\varphi]}$. This set can be computed in polynomial time by Lemma 9.2.

Step 3: Fixing all extensions.

Let $P \in \{0, 1, ?\}^{V \times V}$ and $P' \in \{0, 1, ?\}^{V' \times V'}$ be the ?-block matrices associated with t and t' , respectively. By Lemma 8.1, both P and P' have partition rank at most k . By Theorem 8.2, we can compute canonical extension sets $\text{Ext} \subseteq \{0, 1\}^V$, $\text{Ext}' \subseteq \{0, 1\}^{V'}$ for P, P' of size at most $e(k)$. Recall that for a vertex $v \in \gamma(t)$, by $\text{Ext}(v)$ we denote the set of all extensions of the row \mathbf{p}_v of P in Ext , and similarly for $v' \in V'$ by $\text{Ext}'(v')$ the extensions of $\mathbf{p}'_{v'}$.

If $|\text{Ext}| \neq |\text{Ext}'|$ then $\text{Iso}(G_t, G'_{t'}) = \emptyset$, and we let $\Lambda := \emptyset$. Suppose otherwise. Let $\chi : \text{Ext} \rightarrow \text{Ext}'$ be an arbitrary bijection. We say that an isomorphism ψ from G_t to $G'_{t'}$ agrees with χ if for all $v \in \gamma(t)$ we have $\chi(\text{Ext}(v)) = \text{Ext}'(\psi(v))$ and for every $\mathbf{x} = (x_v \mid v \in V) \in \text{Ext}$ with $\chi(\mathbf{x}) = (x'_{v'} \mid v' \in V')$ and every $v \in \gamma(t)$ it holds that $x_v = x'_{\psi(v)}$.

We argue that every $\psi \in \text{Iso}((G_t)_{D_t, \gamma_t}, (G'_{t'})_{D'_{t'}, \gamma'_{t'}})$ agrees with some bijection $\chi : \text{Ext} \rightarrow \text{Ext}'$. To see this, note first that ψ induces an isomorphism from P to P' . Since Ext and Ext' are canonical, there is a bijection $\chi : \text{Ext} \rightarrow \text{Ext}'$ such that $\chi(\text{Ext}(v)) = \text{Ext}'(\psi(v))$ for all $v \in V$ and for every $\mathbf{x} = (x_v \mid v \in V) \in \text{Ext}$ with $\chi(\mathbf{x}) =: \mathbf{x}' = (x'_{v'} \mid v' \in V') \in \text{Ext}'$ and every $v \in V$ we have $x_v = x'_{\psi(v)}$.

In the following, we fix a bijection $\chi : \text{Ext} \rightarrow \text{Ext}'$, and we let $\text{Iso}((G_t)_{D_t, \gamma_t}, (G'_{t'})_{D'_{t'}, \gamma'_{t'}})_{[\varphi, \chi]}$ be the coset consisting of all $\psi \in \text{Iso}((G_t)_{D_t, \gamma_t}, (G'_{t'})_{D'_{t'}, \gamma'_{t'}})$ that agree with φ and χ . By the usual argument based on Lemma 9.1, it suffices to compute a coset $\Lambda_{\varphi, \chi}$ such that

$$\text{Iso}((G_t)_{D_t, \gamma_t}, (G'_{t'})_{D'_{t'}, \gamma'_{t'}})_{[\varphi, \chi]} \leq \Lambda_{\varphi, \chi} \leq \text{Iso}(G_t, G'_{t'}).$$

Let $u \in N_+^D(t)$ and $u' \in N_+^{D'}(t')$. We say that a $\psi \in \Lambda(u, u')$ agrees with χ if for all $v \in \gamma(u)$ we have $\chi(\text{Ext}(v)) = \text{Ext}'(\psi(v))$ and for every $\mathbf{x} = (x_v \mid v \in V) \in \text{Ext}$ with $\chi(\mathbf{x}) = (x'_{v'} \mid v' \in V')$ and every $v \in \gamma(u)$ it holds that $x_v = x'_{\psi(v)}$. By $\Lambda(u, u')_{[\varphi, \chi]}$ we denote the subcoset of $\Lambda(u, u')$ consisting of all ψ that agree with both φ and χ . We can compute $\Lambda(u, u')_{[\varphi, \chi]}$ by Lemma 9.2.

Let us write $\Lambda(u, u')_{[\varphi, \chi]} = \sigma_{u, u'} \Gamma'_{u, u'}$, where $\Gamma'_{u, u'} \leq \text{Aut}(G'_{u'})$. Observe that

$$g(\mathbf{w}') = \mathbf{w}' \quad \text{for all } g \in \Gamma'_{u, u'} \text{ and } \mathbf{w}' \in W'_{u'}. \quad (9.F)$$

This follows from the fact that all $\psi \in \Lambda(u, u')_{[\varphi, \chi]}$ agree with φ on W_u . Indeed, applying this to $\psi_1 = \sigma_{u, u'}$ and $\psi_2 = \sigma_{u, u'} g^{-1}$, for $\mathbf{w}' \in W'_{u'}$ we have

$$\sigma_{u, u'}^{-1}(\mathbf{w}') = \psi_1^{-1}(\mathbf{w}') = \varphi^{-1}(\mathbf{w}') = \psi_2^{-1}(\mathbf{w}') = \sigma_{u, u'}^{-1}(g(\mathbf{w}'))$$

and thus $\mathbf{w}' = g(\mathbf{w}')$. Furthermore,

$$x'_v = x'_{g(v)} \quad \text{for all } \mathbf{x}' = (x'_v \mid v \in V') \in \text{Ext}' \text{ and all } v \in \gamma'(u'). \quad (9.G)$$

This follows from the fact that all $\psi \in \Lambda(u, u')_{[\varphi, \chi]}$ agree with χ . Let $\psi_1 = \sigma_{u, u'}$ and $\psi_2 = \sigma_{u, u'} g$. Then

$$x'_v = x_{\psi_1^{-1}(v)} = x'_{\psi_2(\psi_1^{-1}(v))} = x'_{g(v)}.$$

For each $g \in \Gamma'_{u, u'}$ we define a permutation \widehat{g} of $V(G'_{u'})$ by

$$\widehat{g}(v) := \begin{cases} g(v) & \text{if } v \in \gamma'(u'), \\ v & \text{if } v \in \gamma'(t') \setminus \gamma'(u'), \\ v & \text{if } v \in W'_{t'}. \end{cases}$$

We let $\widehat{\Gamma}'_{u, u'}$ be the set of all \widehat{g} for $g \in \Gamma'_{u, u'}$. Then $\widehat{\Gamma}'_{u, u'}$ is a subgroup of $\text{Sym}(V(G'_{u'}))$. In fact, the groups $\Gamma'_{u, u'}$ and $\widehat{\Gamma}'_{u, u'}$ are identical as abstract groups, they only differ in their permutation action.

Claim 1. $\widehat{\Gamma}'_{u, u'} \leq \text{Aut}(G'_{t'})$.

Proof. Let $\widehat{g} \in \widehat{\Gamma}'_{u, u'}$ and $v_1, v_2 \in V(G'_t)$ and $v'_i := \widehat{g}(v'_i)$. We shall prove

$$v_1 v_2 \in E(G'_{t'}) \iff v'_1 v'_2 \in E(G'_{t'}). \quad (9.H)$$

Clearly, this equivalence holds if $v_1, v_2 \in \gamma'(u')$, because g is an automorphism of $G'_{u'}$, and also if $v_1, v_2 \notin \gamma'(u')$, because then we have $v_1 = v'_1$ and $v_2 = v'_2$.

So let us assume that $v_1 \in \gamma'(u')$ and $v_2 \notin \gamma'(u')$. Then $v_2 = \widehat{g}(v_2)$. Suppose first that $v_2 = \mathbf{w}' = (w'_v \mid v \in \gamma'(t')) \in W'_{t'}$. Then the u' -projection $\mathbf{w}'|_{u'} = (w'_v \mid v \in \gamma'(u'))$ is in $W'_{u'}$. By (9.F), we have

$$\begin{aligned} v_1 v_2 \in E(G'_{t'}) &\iff w'_{v_1} = 1 \iff v_1 \mathbf{w}'|_{u'} \in E(G'_{u'}) \\ &\iff g(v_1) \mathbf{w}'|_{u'} \in E(G'_{u'}) \iff w'_{g(v_1)} = 1 \iff g(v_1) g(v_2) \in E(G'_{t'}). \end{aligned}$$

It remains to consider the case $v_2 \in \gamma'(t') \setminus \gamma'(u')$. Then there is a $u'' \in N_+^{D'}(t') \setminus \{u'\}$ such that $v_2 \in \gamma'(u'')$. Let $\mathbf{x}' = (x'_v \mid v \in V') \in \text{Ext}'(v_2)$. As v_1 is not contained in $\gamma'(u'')$, the $?$ -index of v_2 , and \mathbf{x}' is an extension of $\mathbf{p}'_{v_2} = (p'_{v_2 v} \mid v \in V')$, we have

$$v_1 v_2 \in E(G'_{t'}) \iff p'_{v_2 v_1} = 1 \iff x'_{v_1} = 1,$$

and similarly $g(v_1) v_2 \in E(G'_{t'}) \iff x'_{g(v_1)} = 1$. By (9.G) and $v_2 = g(v_2)$ we thus have (9.H). \square

Step 4: Admissible bijections.

A bijection $\alpha: N_+^D(t) \rightarrow N_+^{D'}(t')$ is *admissible* if $\Lambda(u, \alpha(u))_{[\varphi, \chi]}$ is nonempty for all $u \in N_+^D(t)$.

Let $\psi \in \text{Iso}(G_t, G'_{t'})$ and $u \in N_+^D(t)$, $u' \in N_+^{D'}(t')$ such that $\psi(\gamma(u)) = \gamma'(u')$. We define the *induced mapping* $\psi|_u: V(G_u) \rightarrow V(G'_{u'})$ by

$$\psi|_u(v) := \begin{cases} \psi(v) & \text{if } v \in \gamma(u), \\ \psi(\mathbf{w})|_{u'} & \text{if } v = \mathbf{w}|_u \text{ for some } \mathbf{w} \in W_t. \end{cases}$$

Claim 2. $\psi|_u \in \text{Iso}(G_u, G'_{u'})$. Furthermore, if ψ agrees with φ and χ , then so does $\psi|_u$.

Proof. We first need to prove that $\psi|_u$ is well-defined. Let $\mathbf{w}_1, \mathbf{w}_2 \in W_t$ such that $\mathbf{w}_1|_u = \mathbf{w}_2|_u$. Suppose that $\mathbf{w}_i = (w_{iv} \mid v \in \gamma(t))$, and let $\mathbf{w}'_i = (w'_{iv'} \mid v' \in \gamma'(t')) := \psi(\mathbf{w}_i)$. Then for all $v \in \gamma(t)$ it holds that

$$w_{iv} = 1 \iff v\mathbf{w}_i \in E(G_t) \iff \psi(v)\mathbf{w}'_i \in E(G'_{t'}) \iff w'_{iv'} = 1.$$

Thus

$$\mathbf{w}_i|_u = (w_{iv} \mid v \in \gamma(u)) = (w'_{iv'} \mid v \in \gamma(u)) = (w'_{iv'} \mid v' \in \gamma'(u')) = \mathbf{w}'_i|_{u'},$$

because $\psi(\gamma(u)) = \gamma'(u')$. As $\mathbf{w}_1|_u = \mathbf{w}_2|_u$, it follows that $\mathbf{w}'_1|_{u'} = \mathbf{w}'_2|_{u'}$. A similar argument shows that $\psi|_u$ is a bijection.

It now follows directly from the fact that ψ is an isomorphism from G_t to $G'_{t'}$ that $\psi|_u$ is an isomorphism, observing that for $v \in \gamma(u)$ and $\mathbf{w} \in W_t$ it holds that $v\mathbf{w} \in E(G_t) \iff v\mathbf{w}|_u \in E(G_u)$.

The second claim follows directly from the definitions of a mapping in $\text{Iso}(G_u, G'_{u'})$ agreeing with φ or χ . \lrcorner

Claim 3. Let $\psi \in \text{Iso}((G_t)_{D_t, \gamma_t}, (G'_{t'})_{D'_{t'}, \gamma'_{t'}})_{[\varphi, \chi]}$. Then there is an admissible bijection α_ψ such that for every $u \in N_+^D(t)$ it holds that $\psi(\gamma(u)) = \gamma'(\alpha_\psi(u))$ and the induced mapping $\psi|_u$ is in $\text{Iso}((G_u)_{D_u, \gamma_u}, (G'_{\alpha_\psi(u)})_{D'_{\alpha_\psi(u)}, \gamma'_{\alpha_\psi(u)}})_{[\varphi, \chi]}$.

Proof. The mapping ψ has an extension $\hat{\psi}$ that is an isomorphism from (D_t, γ_t) to $(D'_{t'}, \gamma'_{t'})$. Let α_ψ be the restriction of $\hat{\psi}$ to $N_+^D(t)$. Then α_ψ is a bijection from $N_+^D(t)$ to $N_+^{D'}(t')$.

Let $u \in N_+^D(t)$. It follows from Claim 2 that the induced mapping $\psi|_u$ is in $\text{Iso}(G_u, G'_{\alpha_\psi(u)})_{[\varphi, \chi]}$. The restriction of $\hat{\psi}$ to D_u is an isomorphism from (D_u, γ_u) to $(D'_{\alpha_\psi(u)}, \gamma'_{\alpha_\psi(u)})$, which shows that $\psi|_u$ respects the decompositions. Hence $\psi|_u \in \text{Iso}((G_u)_{D_u, \gamma_u}, (G'_{\alpha_\psi(u)})_{D'_{\alpha_\psi(u)}, \gamma'_{\alpha_\psi(u)}})_{[\varphi, \chi]}$. \lrcorner

Thus if no admissible bijection exists, then $\text{Iso}((G_t)_{D_t, \gamma_t}, (G'_{t'})_{D'_{t'}, \gamma'_{t'}})_{[\varphi, \chi]} = \emptyset$, and we let $\Lambda_{\varphi, \chi} := \emptyset$. In the following, we assume that there is an admissible bijection. We can find one by computing a perfect matching in the bipartite graph in which u is adjacent to u' if $\Lambda(u, u')_{[\varphi, \chi]}$ is nonempty.

Let α be an admissible bijection. For all $u \in N_+^D(t)$, let $\psi_{u, \alpha(u)} \in \Lambda(u, \alpha(u))_{[\varphi, \chi]}$. We combine the $\psi_{u, \alpha(u)}$ to a mapping ψ_α defined by

$$\psi_\alpha(v) := \begin{cases} \varphi(v) & \text{if } v \in W_t, \\ \psi_{u, \alpha(u)}(v) & \text{if } v \in \gamma(u) \text{ for some } u \in N_+^D(t). \end{cases} \quad (9.I)$$

Observe that for all $u \in N_+^D(t)$ we have $\psi_\alpha|_u = \psi_{u, \alpha(u)}$.

Claim 4. ψ_α is an isomorphism from G_t to $G'_{t'}$.

Proof. As $\gamma(t)$ is the disjoint union of W_t and the $\gamma(u)$ and $\gamma'(t')$ is the disjoint union of $W'_{t'}$ and the $\gamma'(u')$, the mapping ψ_α is well-defined and bijective.

Let $v_1, v_2 \in V(G'_t)$ and $v'_i := \psi_\alpha(v_i)$. We shall prove

$$v_1 v_2 \in E(G_t) \iff v'_1 v'_2 \in E(G'_{t'}). \quad (9.J)$$

Case 1: $v_1, v_2 \in \gamma(u)$ for some $u \in N_+^D(t)$.

Then

$$\begin{aligned} v_1 v_2 \in E(G_t) &\iff v_1 v_2 \in E(G_u) \\ &\iff \psi_{u, \alpha(u)}(v_1) \psi_{u, \alpha(u)}(v_2) \in E(G'_{\alpha(u)}) \iff v'_1 v'_2 \in E(G'_{t'}), \end{aligned}$$

because $\psi_{u, \alpha(u)} \in \Lambda(u, \alpha(u))_{[\varphi, \chi]} \leq \text{Iso}(G_u, G'_{\alpha(u)})$.

Case 2: $v_1 \in \gamma(u)$ for some $u \in N_+^D(t)$ and $v_2 \in W_t$.

Say, $v_2 = \mathbf{w} = (w_v \mid v \in \gamma(t))$ and $v_2' = \varphi(\mathbf{w}) = \mathbf{w}' = (w_{v'}' \mid v' \in \gamma'(t'))$. Then the u -projection $\mathbf{w}|_u$ is in W_u and the u' -projection $\mathbf{w}'|_{\alpha(u)}$ is in $W_{\alpha(u)}'$. As $\psi_{u, \alpha(u)} \in \text{Iso}(G_u, G_{\alpha(u)}')$, we have

$$\begin{aligned} v_1 \mathbf{w} \in E(G_t) &\iff w_{v_1} = 1 \iff v_1 \mathbf{w}|_u \in E(G_u) \\ &\iff v_1' \mathbf{w}'|_{\alpha(u)} \in E(G_{\alpha(u)}') \iff w_{v_1'}' = 1 \iff v_1' \mathbf{w}' \in E(G_{t'}'). \end{aligned}$$

Case 2: $v_2 \in \gamma(u)$ for some $u \in N_+^D(t)$ and $v_1 \in W_t$.

Symmetric to Case 2.

Case 3: $v_1, v_2 \in W_t$.

Then $v_1 v_2 \notin E(G_t)$ and $v_1' v_2' \notin E(G_{t'}')$.

Case 4: $v_1 \in \gamma(u_1)$ and $v_2 \in \gamma(u_2)$ for distinct $u_1, u_2 \in N_+^D(t)$.

As $\psi_{u_2, \alpha(u_2)}$ agrees with χ , we have $\chi(\text{Ext}(v_2)) = \text{Ext}(v_2')$. Let $\mathbf{x} = (x_v \mid v \in V) \in \text{Ext}(v_2)$, and let $\mathbf{x}' = (x_{v'}' \mid v' \in V') := \chi(\mathbf{x})$. Then $\mathbf{x}' \in \text{Ext}(v_2')$. Furthermore, $x_{v_1} = x_{v_1}'$. As $v_1 \notin \gamma(u_2)$, the γ -index of the row \mathbf{p}_{v_2} of P , we have $p_{v_2 v_1} \in \{0, 1\}$, and as \mathbf{x} is an extension of \mathbf{p}_{v_2} , this means that $x_{v_1} = p_{v_2 v_1}$. Similarly, $x_{v_1}' = p_{v_2' v_1}'$. Thus $p_{v_1 v_2} = p_{v_1' v_2}' \in \{0, 1\}$, and this implies (9.J). \lrcorner

Step 5: Closure of the set of admissible bijections.

In this step we prove that without loss of generality we may assume that the sets $\Lambda(u, u')_{[\varphi, \chi]} \neq \emptyset$ have the following closure property.

- (A) For every pair of sequences $u_1, u_2, \dots, u_t \in N_+^D(t)$ and $u_1', u_2', \dots, u_t' \in N_+^{D'}(t')$, if for all $i \in [t]$ we have $\Lambda(u_i, u_i')_{[\varphi, \chi]} \neq \emptyset$ and for all $j \in [t-1]$ we have $\Lambda(u_{j+1}, u_j')_{[\varphi, \chi]} \neq \emptyset$ then $\Lambda(u_0, u_t')_{[\varphi, \chi]} = \emptyset$.

We achieve this as follows. For all $u = u_1, u' = u_t'$ for which there are sequences $u_1, \dots, u_t \in N_+^D(t)$ and $u_1', \dots, u_t' \in N_+^{D'}(t')$ with said properties, we pick for all $i \in [t]$ an element $\nu_i \in \Lambda(u_i, u_i')_{[\varphi, \chi]}$ and for all $j \in [t-1]$ an element $\mu_j \in \Lambda(u_{j+1}, u_j')_{[\varphi, \chi]}$. We then add the map $\nu_1 \mu_1^{-1} \nu_2 \mu_2^{-1} \dots \mu_{t-1}^{-1} \nu_t$ to the set $\Lambda(u_0, u_t')_{[\varphi, \chi]}$. It follows from the fact that isomorphisms compose that this map is an isomorphism in $\text{Iso}(G_{u_0}, G_{u_t'}')_{[\varphi, \chi]}$.

It is also easy to verify that if we perform this procedure once for all pairs u, u' (for which there are sequences with the properties described above) then Property (A) is fulfilled. We thus assume from now on that Property (A) holds. Note that Property (A) implies that for admissible bijections $\alpha_1, \alpha_2, \alpha_3$ the map $\alpha_1 \alpha_2^{-1} \alpha_3$ is also admissible.

Step 6: Construction of the generating set.

We fix an admissible bijection α^0 and for all $u \in N_+^D(t)$, $u' \in N_+^{D'}(t')$ such that $\Lambda(u, u')_{[\varphi, \chi]} \neq \emptyset$, we fix a mapping $\psi_{u, u'}^0 \in \Lambda(u, u')_{[\varphi, \chi]}$. The choice of these mappings is arbitrary and need not be canonical. For every admissible bijection α we define ψ_α^0 as in (9.I) with the map $\psi_{u, u'}^0$ instead of $\psi_{u, u'}$. Furthermore, we let $\psi^0 := \psi_{\alpha^0}^0$.

We say that an admissible bijection α differs from α^0 in at most three positions if there are at most three $u \in N_+^D(t)$ for which $\alpha(u) \neq \alpha^0(u)$. We denote the set of bijections that differ from α^0 in at most three positions by $\alpha^0 \pm 3$. Note that the set $\alpha^0 \pm 3$ can be computed in polynomial time.

Recall the definition of $\widehat{\Gamma}'_{u,u'}$ from Step 3, and let $S_{u,u'}$ be a generating set for $\widehat{\Gamma}'_{u,u'}$ (the succinct representation of $\Lambda(u, u')_{[\varphi, \chi]}$ contains such a generating set). We let

$$S := \{(\psi_\alpha^0)^{-1}\psi_{\alpha'}^0 \mid \alpha, \alpha' \in (\alpha^0 \pm 3)\} \cup \bigcup_{u, u'} S_{u, u'},$$

and we let

$$\Lambda_{\varphi, \chi} := \psi^0 \langle S \rangle.$$

Note that $\psi^0 \in \Lambda_{\varphi, \chi}$ and $\psi_\alpha^0 = \psi^0(\psi^0)^{-1}\psi_\alpha^0 \in \Lambda_{\varphi, \chi}$ for all $\alpha \in \alpha^0 \pm 3$.

Claim 5. $\text{Iso}((G_t)_{D_t, \gamma_t}, (G'_{t'})_{D'_{t'}, \gamma'_{t'}})_{[\varphi, \chi]} \leq \Lambda_{\varphi, \chi} \leq \text{Iso}(G_t, G'_{t'})$.

Proof. It follows from Claim 4 that $\psi^0(\psi_\alpha^0)^{-1}\psi_{\alpha'}^0 \in \text{Iso}(G_t, G'_{t'})$ for all admissible bijections α, α' , and it follows from Claim 1 that $\psi^0 g \in \text{Iso}(G_t, G'_{t'})$ for all $g \in S_{u, u'} \subseteq \widehat{\Gamma}'_{u, u'}$. This proves $\Lambda_{\varphi, \chi} \leq \text{Iso}(G_t, G'_{t'})$.

To prove $\text{Iso}((G_t)_{D_t, \gamma_t}, (G'_{t'})_{D'_{t'}, \gamma'_{t'}})_{[\varphi, \chi]} \leq \Lambda_{\varphi, \chi}$, we consider the subset K of $\text{Iso}(G_t, G'_{t'})_{[\varphi, \chi]}$ consisting of all ψ satisfying the following two conditions:

- (i) There is an admissible bijection α_ψ such that $\psi(\gamma(u)) = \gamma'(\alpha_\psi(u))$ for all $u \in N_+^D(t)$.
- (ii) $\psi|_u \in \Lambda(u, \alpha_\psi(u))_{[\varphi, \chi]}$ for all $u \in N_+^D(t)$.

It follows from Claim 3 that every $\psi \in \text{Iso}((G_t)_{D_t, \gamma_t}, (G'_{t'})_{D'_{t'}, \gamma'_{t'}})_{[\varphi, \chi]}$ is contained in K . Furthermore, for every admissible bijection α we have $\psi_\alpha^0 \in K$, because $\psi_\alpha^0|_u = \psi_{u, \alpha(u)}^0 \in \Lambda(u, \alpha_\psi(u))_{[\varphi, \chi]}$ for all $u \in N_+^D(t)$. In particular, $\psi^0 = \psi_{\alpha^0}^0 \in K$.

We shall prove that

$$K \subseteq \Lambda_{\varphi, \chi}.$$

For $\psi \in K$, we prove

$$\psi \in \Lambda_{\varphi, \chi}.$$

The map $\psi_{\alpha_\psi}^0$ agrees with ψ up to elements in $\widehat{\Gamma}'_{u, \alpha_\psi(u)}$. More precisely, $\psi|_u \in \Lambda(u, \alpha_\psi(u))_{[\varphi, \chi]}$ can be written as $\psi|_u = \psi_{u, \alpha_\psi(u)}^0 g_u$ for some $g_u \in \Gamma'_{u, \alpha_\psi(u)}$. As the cones $\gamma(u')$ for $u' \in N_+^D(t')$ are mutually disjoint, the permutations $\widehat{g}_u \in \widehat{\Gamma}'_{u, \alpha(u)}$ (see Step 3) commute. Thus we have

$$\psi = \psi_{\alpha_\psi}^0 g_{u_1} \cdots g_{u_m}$$

for an arbitrary enumeration u_1, \dots, u_m of $N_+^D(t)$. This proves that

$$\psi_{\alpha_\psi}^0 \in \Lambda_{\varphi, \chi} \iff \psi \in \Lambda_{\varphi, \chi}.$$

It thus suffices to show that $\psi_{\alpha_\psi}^0 \in \Lambda_{\varphi, \chi}$. More generally, we now show for an arbitrary admissible bijection α^* that $\psi_{\alpha^*}^0 \in \Lambda_{\varphi, \chi}$. We show this by induction on the number $d(\alpha^*)$ of $u \in N_+^D(t)$ such that $\alpha^*(u) \neq \alpha^0(u)$.

In the base case $d(\alpha^*) = 0$ we have $\alpha^* = \alpha^0$ and thus $\psi_{\alpha^*}^0 = \psi^0 \in \Lambda_{\varphi, \chi}$.

For the inductive step, suppose that $\alpha^* \neq \alpha^0$. Then there are at least two $u \in N_+^D(t)$ such that $\alpha^*(u) \neq \alpha(u)$.

Case 1: There are distinct $u_1, u_2 \in N_+^D(t)$ such that $\alpha^*(u_1) = \alpha^0(u_2)$ and $\alpha^*(u_2) = \alpha^0(u_1)$.

Let α be the bijection from $N_+^D(t)$ to $N_+^{D'}(t')$ with $\alpha(u_i) = \alpha^*(u_i)$ for $i = 1, 2$ and $\alpha(u) = \alpha^0(u)$ for all $u \neq u_1, u_2$. Then α is admissible, because $\alpha(u_i) = \alpha^*(u_i)$ for $i = 1, 2$ and $\alpha(u) = \alpha^0(u)$ for all $u \neq u_1, u_2$. Hence $\alpha \in \alpha^0 \pm 3$.

Let $\alpha' = \alpha^* \alpha^{-1} \alpha^0$, which is admissible by Property (A). Then $d(\alpha') = d(\alpha^*) - 2$. Let $\psi' := \psi_{\alpha^*}^0 (\psi_\alpha^0)^{-1} \psi^0$. Then $\psi' \in K$ with $\alpha_{\psi'} = \alpha'$. By the inductive hypothesis, $\psi' \in \Lambda_{\varphi, \chi}$. Furthermore, $\psi^0 \in \Lambda_{\varphi, \chi}$ and $\psi_\alpha^0 \in \Lambda_{\varphi, \chi}$, the latter because $\alpha \in \alpha^0 \pm 3$. Thus $\psi_{\alpha^*}^0 = \psi'(\psi^0)^{-1}\psi_\alpha^0 \in \Lambda_{\varphi, \chi}$.

Case 2: There are distinct $u_1, u_2, u_3 \in N_+^D(t)$ such that $\alpha^*(u_1) = \alpha^0(u_2)$ and $\alpha^*(u_2) = \alpha^0(u_3)$.

Let α_1, α_2 be the bijections from $N_+^D(t)$ to $N_+^{D'}(t')$ defined by

$$\begin{aligned} \alpha_1(u_1) &:= \alpha^0(u_2), \\ \alpha_1(u_2) &:= \alpha^0(u_3), \\ \alpha_1(u_3) &:= \alpha^0(u_1), \\ \alpha_1(u) &:= \alpha^0(u) && \text{for all } u \in N_+^D(t) \setminus \{u_1, u_2, u_3\}, \\ \alpha_2(u_1) &:= \alpha^0(u_3), \\ \alpha_2(u_2) &:= \alpha^0(u_2), \\ \alpha_2(u_3) &:= \alpha^0(u_1), \\ \alpha_2(u) &:= \alpha^0(u) && \text{for all } u \in N_+^D(t) \setminus \{u_1, u_2, u_3\}. \end{aligned}$$

To prove that α_1 is admissible, we need to prove that $\Lambda(u, \alpha_1(u))_{[\varphi, \chi]} \neq \emptyset$ for all $u \in N_+^D(t)$. For $u \neq u_3$ this is obvious since for such a u we have $\alpha_1(u) = \alpha^0(u)$ or $\alpha_1(u) = \alpha^*(u)$ and both α^0 and α^* are admissible. For $u = u_3$ we observe that the two sequences u_3, u_2, u_1 and $\alpha^0(u_3), \alpha^0(u_2), \alpha^0(u_1)$ satisfy the assumptions on the sequences in Property(A) and we therefore conclude $\Lambda(u_3, \alpha^0(u_1))_{[\varphi, \chi]} \neq \emptyset$.

The proof that α_2 is admissible is similar.

Since both α_1 and α_2 differ from α^0 in at most 3 places, we have $\alpha_1, \alpha_2 \in \alpha^0 \pm 3$.

Let $\alpha' = \alpha^* \alpha_1^{-1} \alpha_2$, which is admissible by Property (A). Then $\alpha'(u_1) = \alpha^0(u_3)$, $\alpha'(u_2) = \alpha^0(u_2)$, and $\alpha'(u) = \alpha^*(u)$ for all $u \notin \{u_1, u_2\}$. Thus $d(\alpha') = d(\alpha^*) - 1$. Let $\psi' := \psi_{\alpha^*}^0 (\psi_{\alpha_1}^0)^{-1} \psi_{\alpha_2}^0$. Then $\alpha_{\psi'} = \alpha'$. Furthermore, $\psi_{\alpha_1}^0, \psi_{\alpha_2}^0 \in \Lambda_{\varphi, \chi}$. Thus it suffices now show that $\psi' \in K$, since this implies $\psi' \in \Lambda_{\varphi, \chi}$ by the inductive hypothesis, which implies $\psi_{\alpha^*}^0 = \psi' (\psi_{\alpha_2}^0)^{-1} \psi_{\alpha_1}^0 \in \Lambda_{\varphi, \psi}$.

To see that $\psi' \in K$ we need to show that $\psi'|_u \in \Lambda(u, \alpha^*(u))_{[\varphi, \chi]}$ for all $u \in N_+^D(t)$.

Note that by the definition of α_1 and α_2 , for all $u' \in N_+^{D'}(t') \setminus \{\alpha^*(u_1), \alpha^*(u_2)\}$ we have $\alpha_1^{-1}(u') = \alpha_2^{-1}(u')$.

Thus if $u \in N_+^D(t) \setminus \{u_1, u_2\}$ then

$$\psi'|_u = \psi_{\alpha^*}^0|_u (\psi_{\alpha_1}^0|_{\alpha_1^{-1}(\alpha^*(u))})^{-1} \psi_{\alpha_2}^0|_{\alpha_2^{-1}(\alpha^*(u))} = \psi_{\alpha^*}^0|_u \in \Lambda(u, \alpha^*(u))_{[\varphi, \chi]}.$$

On the other hand if $u \in \{u_1, u_2\}$ then $\psi_{\alpha^*}^0|_u = \psi_{\alpha_1}^0|_u$ and thus

$$\psi'|_u = \psi_{\alpha^*}^0|_u (\psi_{\alpha_1}^0|_u)^{-1} \psi_{\alpha_2}^0|_u = \psi_{\alpha_2}^0|_u \in \Lambda(u, \alpha_2(u))_{[\varphi, \chi]}.$$

This shows that $\psi_{\alpha^*}^0 \in \Lambda_{\varphi, \chi}$ for all admissible α^* and thus in particular $\psi_{\alpha_\psi}^0 \in \Lambda_{\varphi, \chi}$, finishing the proof of the claim. \lrcorner

Since $\Lambda_{\varphi, \chi}$ can be computed in polynomial time for all choices of φ and χ the theorem follows. \square

Proof of Theorem 9.3. Using dynamic programming and the previous lemma, we can compute for all $t \in D$ and $t' \in D$ a coset $\Lambda(t, t')$ satisfying $\text{Iso}((G_t)_{D_t, \gamma_t}, (G_{t'})_{D_{t'}, \gamma_{t'}}) \subseteq \Lambda(t, t') \subseteq \text{Iso}(G_t, G_{t'})$. We let $\Lambda := \Lambda_{r, r'}$ for the roots r, r' of D, D' , respectively. \square

Theorem 9.5. *For every $k \in \mathbb{N}$ there is a polynomial time algorithm that, given graphs G and G' of rank width at most k , computes the set $\text{Iso}(G, G')$ of all isomorphisms from G to G' .*

Proof. By Theorem 7.2 we can compute for G and G' canonical treelike decompositions (D, γ) and (D', γ') of width at most $a(k)$. By Lemma 4.1 we can assume that these decompositions are normal. Since these decompositions are canonical, $\text{Iso}((G)_{D, \gamma}, (G')_{D', \gamma'}) = \text{Iso}(G, G')$. The theorem now follows directly from Theorem 9.3. \square

10 Conclusions

For every fixed k we obtain a polynomial time isomorphism test for graph classes of bounded rank width, unfortunately with a horrible running time: we only have a non-elementary upper bound (in terms of k) for the degree of the polynomial bounding the running time. Thus before even asking whether the isomorphism problem is fixed-parameter tractable if parameterized by rank-width, we ask for an algorithm with a running time $n^{O(k)}$. The bottleneck is the bound we obtain for the size of a triple cover of a tangle (see Lemma 3.1); our algorithm has to enumerate all triple covers of all maximal tangles. But maybe there is a way to avoid this.

Our algorithm uses the group theoretic machinery, but the group theory involved is fairly elementary. It seems conceivable that it can be avoided altogether and there is a combinatorial algorithm deciding isomorphism of rank width at most k . Specifically, we ask whether for any k there is an ℓ such that the ℓ -dimensional Weisfeiler-Lehman algorithm decides isomorphism of graphs of rank width at most k .

Most of the arguments that we use in the construction of canonical bounded width decompositions apply to arbitrary connectivity functions and not just the cut rank function. (Only from Section 6 onwards we use specific properties of the cut rank function.) It is an interesting question whether there is a polynomial time isomorphism test for arbitrary connectivity functions of bounded branch width. Even if this is not the case, it would be interesting to understand for which connectivity functions beyond the cut rank function such an isomorphism test exists.

In the end, the main question is whether our results help to solve the isomorphism problem for general graphs. The immediate answer is ‘no’. However, we do believe that structural techniques such as those developed here (and also in [12, 14]), in combination with group theoretic techniques, may help to design graph isomorphism test with an improved worst-case running time.

References

- [1] L. Babai, D.Yu. Grigoryev, and D.M. Mount. Isomorphism of graphs with bounded eigenvalue multiplicity. In *Proceedings of the 14th ACM Symposium on Theory of Computing*, pages 310–324, 1982.
- [2] H.L. Bodlaender. Polynomial algorithms for graph isomorphism and chromatic index on partial k -trees. *Journal of Algorithms*, 11:631–643, 1990.
- [3] Kellogg S. Booth and C. J. Colbourn. Problems polynomially equivalent to graph isomorphism. Technical Report CS-77-04, Comp. Sci. Dep., Univ. Waterloo, 1979.
- [4] D. G. Corneil, H. Lerchs, and L. Stewart Burlingham. Complement reducible graphs. *Discrete Applied Mathematics*, 3(3):163–174, 1981.
- [5] B. Courcelle, J.A. Makowsky, and U. Rotics. On the fixed-parameter complexity of graph enumeration problems definable in monadic second-order logic. *Discrete Applied Mathematics*, 108(1–2):23–52, 2001.
- [6] B. Courcelle and S. Olariu. Upper bounds to the clique-width of graphs. *Discrete Applied Mathematics*, 101:77–114, 2000.
- [7] Andrew Curtis, Min Lin, Ross McConnell, Yahav Nussbaum, Francisco Soullignac, Jeremy Spinrad, and Jayme Szwarcfiter. Isomorphism of graph classes related to the circular-ones property. *Discrete Mathematics and Theoretical Computer Science*, 15(1):157–182, 2013.
- [8] W. Espelage, F. Gurski, and E. Wanke. How to solve NP-hard graph problems on clique-width bounded graphs in polynomial time. In A. Brandstädt and V. Le, editors, *Proceedings*

- of the 27th Workshop on Graph-Theoretic Concepts in Computer Science, volume 2204 of *Lecture Notes in Computer Science*, pages 117–128. Springer-Verlag, 2001.
- [9] I. S. Filotti and J. N. Mayer. A polynomial-time algorithm for determining the isomorphism of graphs of fixed genus. In *Proceedings of the 12th ACM Symposium on Theory of Computing*, pages 236–243, 1980.
 - [10] E. Fischer, J.A. Makowsky, and E.V. Ravve. Counting truth assignments of formulas of bounded tree-width or clique-width. *Discrete Applied Mathematics*, 156(4):511 – 529, 2008.
 - [11] J. Geelen, B. Gerards, and G. Whittle. Tangles, tree-decompositions and grids in matroids. *Journal of Combinatorial Theory, Series B*, 99(4):657–667, 2009.
 - [12] M. Grohe. Descriptive complexity, canonisation, and definable graph structure theory. Manuscript available at <http://www.lii.rwth-aachen.de/de/mitarbeiter/13-mitarbeiter/professoren/39-book-descrip>
 - [13] M. Grohe. Definable tree decompositions. In *Proceedings of the 23rd IEEE Symposium on Logic in Computer Science*, pages 406–417, 2008.
 - [14] M. Grohe and D. Marx. Structure theorem and isomorphism test for graphs with excluded topological subgraphs. *SIAM Journal on Computing*, 44(1):114–159, 2015.
 - [15] M. Grohe and P. Schweitzer. Computing with tangles. *ArXiv*, arXiv:1503.00190 [cs.DM], 2015. Conference version to appear in STOC’15 Proceedings.
 - [16] J.E. Hopcroft and R. Tarjan. Isomorphism of planar graphs (working paper). In R. E. Miller and J. W. Thatcher, editors, *Complexity of Computer Computations*. Plenum Press, 1972.
 - [17] S. Iwata, L. Fleischer, and S. Fujishige. A combinatorial strongly polynomial algorithm for minimizing submodular functions. *Journal of the ACM*, 48(4):761–777, 2001.
 - [18] D. Kobler and U. Rotics. Edge dominating set and colorings on graphs with fixed clique-width. *Discrete Applied Mathematics*, 126(2–3):197 – 221, 2003.
 - [19] Johannes Köbler, Sebastian Kuhnert, and Oleg Verbitsky. Helly circular-arc graph isomorphism is in logspace. In Krishnendu Chatterjee and Jiri Sgall, editors, *Mathematical Foundations of Computer Science 2013 - 38th International Symposium*, volume 8087 of *Lecture Notes in Computer Science*, pages 631–642. Springer, 2013.
 - [20] Stefan Kratsch and Pascal Schweitzer. Graph isomorphism for graph classes characterized by two forbidden induced subgraphs. In *Graph-Theoretic Concepts in Computer Science - 38th International Workshop*, volume 7551 of *Lecture Notes in Computer Science*, pages 34–45. Springer, 2012.
 - [21] Vincent Limouzy, Fabien de Montgolfier, and Michaël Rao. NLC-2 graph recognition and isomorphism. In Andreas Brandstädt, Dieter Kratsch, and Haiko Müller, editors, *Graph-Theoretic Concepts in Computer Science, 33rd International Workshop, WG 2007, Dornburg, Germany, June 21-23, 2007. Revised Papers*, volume 4769 of *Lecture Notes in Computer Science*, pages 86–98. Springer, 2007.
 - [22] D. Lokshtanov, M. Pilipczuk, M. Pilipczuk, and S. Saurabh. Fixed-parameter tractable canonization and isomorphism test for graphs of bounded treewidth. In *Proceedings of the 55th Annual IEEE Symposium on Foundations of Computer Science*, pages 186–195, 2014.

- [23] G.S. Lueker and K.S. Booth. A linear time algorithm for deciding interval graph isomorphism. *Journal of the ACM*, 26(2):183–195, 1979.
- [24] E.M. Luks. Isomorphism of graphs of bounded valance can be tested in polynomial time. *Journal of Computer and System Sciences*, 25:42–65, 1982.
- [25] G. L. Miller. Isomorphism testing for graphs of bounded genus. In *Proceedings of the 12th ACM Symposium on Theory of Computing*, pages 225–235, 1980.
- [26] S.-I. Oum and P.D. Seymour. Approximating clique-width and branch-width. *Journal of Combinatorial Theory, Series B*, 96:514–528, 2006.
- [27] J. Oxley. *Matroid Theory*. Cambridge University Press, 2nd edition, 2011.
- [28] I. N. Ponomarenko. The isomorphism problem for classes of graphs that are invariant with respect to contraction. *Zap. Nauchn. Sem. Leningrad. Otdel. Mat. Inst. Steklov. (LOMI)*, 174(Teor. Slozhn. Vychisl. 3):147–177, 182, 1988. In Russian.
- [29] N. Robertson and P.D. Seymour. Graph minors X. Obstructions to tree-decomposition. *Journal of Combinatorial Theory, Series B*, 52:153–190, 1991.
- [30] S.-I-Oum and P. Seymour. Testing branch-width. *Journal of Combinatorial Theory, Series B*, 97:385–393, 2007.
- [31] A. Schrijver. A combinatorial algorithm minimizing submodular functions in strongly polynomial time. *Journal of Combinatorial Theory, Series B*, 80(2):346–355, 2000.
- [32] Pascal Schweitzer. Towards an isomorphism dichotomy for hereditary graph classes. In *32nd International Symposium on Theoretical Aspects of Computer Science*, volume 30 of *LIPICs*, pages 689–702, 2015.
- [33] Ákos Seress. *Permutation Group Algorithms*. Cambridge Tracts in Mathematics. Cambridge University Press, 2003.